


Wireless Networking Survival Guide



TECHREPUBLIC RESOURCE CD LICENSE AGREEMENT

READ THIS AGREEMENT BEFORE USING THIS TECHREPUBLIC RESOURCE CD-ROM DISK ("CD") FROM TECHREPUBLIC. BY USING THE CD YOU AGREE TO BE BOUND BY THE TERMS AND CONDITIONS OF THIS AGREEMENT. IF YOU DO NOT AGREE TO THE TERMS AND CONDITIONS OF THIS AGREEMENT, IMMEDIATELY RETURN THE UNUSED CD FOR A FULL REFUND OF MONIES PAID, IF ANY.

The articles, forms, tools, templates, programs, and other materials included on this CD and their compilation (the "Collection") are licensed to you subject to the terms and conditions of this Agreement by TechRepublic, having a place of business at 1630 Lyndon Farm Ct, Louisville, KY 40223 ("TechRepublic"). By using the Collection, in whole or in part, you agree to be bound by the terms and conditions of this Agreement. TechRepublic owns the title to the Collection and to all intellectual property rights therein, except in so far as it contains materials that are proprietary to third-party suppliers. All rights in the Collection except those expressly granted to you in this Agreement are reserved to TechRepublic and such suppliers, as their respective interests may appear.

1. Limited License

TechRepublic grants you a limited, nonexclusive, nontransferable license to use the Collection on a single dedicated computer. This Agreement and your rights hereunder shall automatically terminate if you fail to comply with any provision of this Agreement. Upon such termination, you agree to destroy the CD and all copies of the CD, whether or not lawful, that are in your possession or under your control.

2. Additional Restrictions

A. You shall not (and shall not permit other persons or entities to) directly or indirectly, by electronic or other means, copy or reproduce (except for archival purposes as permitted by law), publish, distribute, rent, lease, sell, sublicense, assign, or otherwise transfer the Collection or any part thereof or this Agreement, and neither the CD nor its contents can be shared over a network for access by multiple users without a separate site license agreement. Any attempt to do so shall be void and of no effect.

B. You shall not (and shall not permit other persons or entities to) reverse-engineer, decompile, disassemble, merge, modify, create derivative works of, or translate the Collection or use the Collection for any purpose.

C. You shall not (and shall not permit other persons or entities to) remove or obscure TechRepublic's or its suppliers' copyright, trademark, or other proprietary notices or legends from any portion of the Collection or any related materials.

3. Limited Warranty and Limited Liability

A. THE ONLY WARRANTY MADE BY TECHREPUBLIC IS THAT THE ORIGINAL CD IN WHICH THE COLLECTION IS EMBODIED AND WHICH IS DISTRIBUTED BY TECHREPUBLIC SHALL BE FREE OF DEFECTS IN MATERIALS AND WORKMANSHIP FOR A PERIOD OF NINETY (90) DAYS AFTER DELIVERY TO YOU. TECHREPUBLIC'S AND ITS SUPPLIERS' ENTIRE LIABILITY AND YOUR EXCLUSIVE REMEDY SHALL BE LIMITED TO THE REPLACEMENT OF THE ORIGINAL CD, IF DEFECTIVE, WITHIN A REASONABLE PERIOD OF TIME.

B. EXCEPT AS SPECIFICALLY PROVIDED ABOVE, THE COLLECTION IS PROVIDED 'AS IS' WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, ANY WARRANTY OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE RESULTS AND PERFORMANCE OF THE SOFTWARE AND OTHER MATERIAL THAT IS PART OF THE COLLECTION IS ASSUMED BY YOU. AND TECHREPUBLIC AND ITS SUPPLIERS ASSUME NO RESPONSIBILITY FOR THE ACCURACY ON APPLICATION OF OR ERRORS OR OMISSIONS IN THE COLLECTION. IN NO EVENT SHALL TECHREPUBLIC OR ITS SUPPLIERS BE LIABLE FOR ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE COLLECTION, EVEN IF TECHREPUBLIC OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE LIKELIHOOD OF SUCH DAMAGES OCCURRING. TECHREPUBLIC AND ITS SUPPLIERS SHALL NOT BE LIABLE FOR ANY LOSS, DAMAGES, OR COSTS ARISING OUT OF, BUT NOT LIMITED TO, LOST PROFITS OR REVENUE; LOSS OF USE OF THE COLLECTION; LOSS OF DATA OR EQUIPMENT;

COST OF RECOVERING SOFTWARE, DATA, OR THE MATERIALS IN THE COLLECTION; THE COST OF SUBSTITUTE SOFTWARE, DATA OR MATERIALS IN THE COLLECTION; CLAIMS BY THIRD PARTIES; OR OTHER SIMILAR COSTS.

C. THE WARRANTIES AND REMEDIES SET FORTH HEREIN ARE EXCLUSIVE AND IN LIEU OF ALL OTHERS, ORAL OR WRITTEN, EXPRESSED OR IMPLIED. NO TECHREPUBLIC AGENT OR EMPLOYEE OR THIRD PARTY IS AUTHORIZED TO MAKE ANY MODIFICATION OR ADDITION TO THIS WARRANTY.

D. SOME STATES DO NOT ALLOW EXCLUSION OR LIMITATION OF IMPLIED WARRANTIES OR LIMITATION OF LIABILITY FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES, SO THE ABOVE LIMITATIONS OR EXCLUSIONS MAY NOT APPLY TO YOU.

4. U.S. Government Restricted Rights

The Collection is licensed subject to RESTRICTED RIGHTS. Use, duplication, or disclosure by the U.S. Government or any person or entity acting on its behalf is subject to restrictions as set forth in subdivision (c)(1)(iii) of the Rights in Technical Data and Computer Software Clause at DFARS (48 CFR 252.227-7013) for DoD contracts, in paragraphs (c)(1) and (2) of the Commercial Computer Software and the Restricted Rights clause in the FAR (48 CFR 52.227-19) for civilian agencies or in other comparable agency clauses. The contractor, manufacturer, is TechRepublic.

5. General Provision

Nothing in this Agreement constitutes a waiver of TechRepublic's or its suppliers' rights under U.S. copyright laws or any other federal, state, local, or foreign law. You are responsible for installation, management, and operation of the Collection. This Agreement shall be construed, interpreted, and governed under California law.

CD-ROM Requirements

The TechRepublic Resource CD requires:

- Windows 98/98SE/ME/NT4/2000 or XP
- Internet Explorer 5.0 or later
- 16 MB of RAM or more
- 10 MB of free disk space or more
- Windows-compatible CD-ROM drive

Wireless Networking Survival Guide

Copyright

©1995-2003 by CNET Networks, Inc. All rights reserved. TechRepublic and its logo are trademarks of CNET Networks, Inc. All other product names or services identified throughout this book are trademarks or registered trademarks of their respective companies. Reproduction of this publication in any form without prior written permission is forbidden.

Disclaimer

The information contained herein has been obtained from sources believed to be reliable. CNET Networks, Inc. disclaims all warranties as to the accuracy, completeness, or adequacy of such information. CNET Networks, Inc. shall have no liability for errors, omissions, or inadequacies in the information contained herein or for the interpretations thereof. The reader assumes sole responsibility for the selection of these materials to achieve its intended results. The opinions expressed herein are subject to change without notice.

CD-ROM License

TechRepublic grants you a limited, nonexclusive, nontransferable license to use the CD-ROM on a single dedicated computer. The use of the TechRepublic Resource CD is governed by the license agreement that can be found in the printed documentation included with the CD-ROM. Read the agreement carefully before using the CD-ROM that accompanies this book.

Contact Us

TechRepublic

1630 Lyndon Farm Court

Louisville, KY 40223

E-mail: customerservice@techrepublic.com

Tel.: 1.800.217.4339

www.techrepublic.com

ISBN 1-932509-01-1

October 2003

B059

Credits

Vice President, TechRepublic

Bob Artner

Assistant Vice President, TechRepublic

Kimberly Henderson

Executive Editor, Premium Products

Erik Eckel

Managing Editor, Premium Products

Janice Conard

Content Resources Manager

Marilyn Bryan

Graphic Artists

Natalie Eckerle

Kimberly Wright

Executive Editor, TechRepublic and Builder.com

Veronica Combs

Senior Editors

Paul Baldwin

Beth Blakely

Toni Bowers

Bill Detwiler

Jason Hiner

Judy Mottl

John Sheesley

Jim Wells

Review Edit Manager

Rich Crossett

Review Editors

Kachina Dunn

Jody Gilbert

Kim Mays

Amy Sellers

Copy Editors

Selena Frye

Joyce Mathai

Suzanne Thornberry

Julie Tonini

Linda Watkins

Membership Director

Dan Scofield

Promotions Manager

Megan Hancock

Foreword

Wireless networking innovations enable a host of advantages. Antiquated tethers no longer restrict desktops, laptops, and handheld computers. Instead, systems can be placed where most convenient, even if cable runs aren't readily available. Best of all, mobile systems become truly mobile platforms.

Wireless e-mail access, Internet, intranet, and extranet usage, file exchanges, and other forms of cordless collaboration are now routine. Wireless devices provide additional opportunities when rolling out new systems and rearranging older configurations, too. This is true whether you're working from a family room, small office, cubicle, or server room.

For all of their benefits, however, many issues still plague wireless network configuration and administration. Implementing and maintaining a secure and efficient wireless network requires careful planning and diligent administration. You must take care to ensure that 802.11 networks are properly configured to enable wireless access while guarding against opening your network and all of its data to unauthorized use.

TechRepublic's *Wireless Networking Survival Guide* reviews wireless networking fundamentals, describes important configuration and troubleshooting techniques, lists critical security precautions, and includes helpful information on popular products and devices. The *Wireless Networking Survival Guide* book also includes a companion CD-ROM, the *Wireless Networking Tool Kit*, which collects helpful templates, diagrams, and checklists you can use to ease wireless network administration.

This unique book and CD-ROM set won't sit neglected on your bookshelf. Our editors and real-world IT professionals have created these aids to be resources you'll rely upon and utilize regularly. You'll find that the spiral-bound book stays open where you need it and doesn't flip closed the moment you set it down. The CD-ROM includes ready-to-use, customizable charts and templates, ensuring you maximize the tools' benefits.

With the *Wireless Networking Survival Guide* and *Wireless Networking Tool Kit*, you'll have all you need to:

- ▶ Understand wireless network operation.
- ▶ Add and configure network adapters.
- ▶ Configure wireless network connections.
- ▶ Establish file access permissions.
- ▶ Secure wireless systems and networks.
- ▶ Select the best products for your needs.
- ▶ Roll out and maintain a local area network and wireless connections.

Learn from professionals in the field. Take advantage of TechRepublic's proven solutions to ensure that your wireless network operates as efficiently and securely as possible.

If you have suggestions or comments regarding this TechRepublic product, please e-mail us at trproducts@techrepublic.com.

Quick Reference

Wireless Fundamentals	1
Configuration and Troubleshooting	21
File and Share Permissions	73
Wireless Security.....	95
Products and Reviews	129

Wireless Networking Survival Guide

Wireless Fundamentals

Wireless—The real thing, finally	1
Look Ma, no wires!.....	2
Use wireless technology to triumph over networking nightmares	5
Take advantage of the cost savings of a wireless LAN.....	6
Get the scoop on WLANs with this wireless networking overview	8
A primer on Wireless Application Protocol (WAP)	10
Understanding wireless LAN protocols and components.....	12
Evaluating the wireless networking options.....	15

Configuration and Troubleshooting

Bridge floors and buildings with wireless access points	21
Span the WAN with wireless bridges	23
Plan effectively and save big on a wireless bridging deployment	27
Add and configure network adapters	30
Add protocols, services, and network clients and bind them all to your NIC.....	35
Understanding wireless network settings	40
Windows XP offers groundbreaking WLAN functionality.....	44
Configuring a wireless LAN connection in Windows XP.....	46
Create local user accounts for Windows 2K/XP peer-to-peer networking	49
Install a wireless connection on your home network.....	53
Diagnosing wireless network performance problems	57
Fix hardware and configuration issues common to wireless LANs.....	61
Troubleshooting the wireless woes	64
Troubleshoot wireless networking antennas	66

File and Share Permissions

File-sharing permissions in Windows 2000	73
NTFS permissions in Windows 2000	77
Combining sharing and NTFS permissions in Windows 2000	80
Establish the correct file-sharing permissions in Windows XP	83
Effectively set and troubleshoot NTFS permissions in Windows XP	86
Combining sharing and NTFS permissions in Windows XP	90

Wireless Security

Keep up with public wireless dangers and Wi-Fi security standards.....	95
Design a secure wireless LAN	97
Think security when setting up an 802.11b wireless network.....	100
How to beef up wireless security	102
Use WEP to improve security on your wireless network	105
Take steps to secure vulnerable WLANs	110
At last, real wireless LAN security.....	111
WPA wireless security offers multiple advantages over WEP	113
Six tips for implementing closed networking on a wireless network.....	115
Don't use MAC filtering as your only wireless network security solution	116
Choosing a vendor solution for wireless LAN security with 802.1x and EAP	119
Follow these steps to tighten security on Linksys wireless networks	121
XP client configuration for enhanced security on a Linksys wireless network	124

Products and Reviews

How to select the right wireless hardware for your home network.....	129
Go wireless with 802.11 options from Dell and Gateway	131
Supporting wireless users with 802.11 options from Compaq and IBM	132
Cut the cord with Agere Wireless USB Client systems	134
ORiNOCO's wireless network: Avoid its sticky setup problems	139
Installing ORiNOCO wireless gateway is a snap	142
ORiNOCO USB client setup makes a turn for the better	144
3Com AirConnect: Wireless for the great wide open	147

A review of 3Com's HomeConnect Home Wireless Gateway	154
Product Rating: 3Com Home Wireless Gateway	158
Connect wires and wireless with the Linksys Ethernet Bridge	160
Product Rating: Linksys EtherFast wireless AP and cable/DSL router with 4-port	164
Product Rating: NetGear MR314 cable/DSL wireless router	166
Product Rating: HP wireless gateway hn200w	167
Product Rating: Intel AnyPoint wireless gateway	169
Product Rating: SMC Barricade wireless broadband router	170
Product Rating: SMC EZ Connect 802.11a wireless access point.....	172
Quickly add wireless ports with SMC's EZ Connect wireless access point	174
Untether your network with SMC's wireless adapter	176
SMC's wireless broadband router offers performance tempered with caveats	178
Vivato's WLAN switches extend Wi-Fi range	182

Wireless Fundamentals

Wireless—The real thing, finally	1
Look Ma, no wires!	2
Use wireless technology to triumph over networking nightmares	5
Take advantage of the cost savings of a wireless LAN	6
Get the scoop on WLANs with this wireless networking overview	8
A primer on Wireless Application Protocol (WAP)	10
Understanding wireless LAN protocols and components	12
Evaluating the wireless networking options	15

Wireless—The real thing, finally

Jun 7, 2002

By Rich Castagna, ZDNet

The computer business is famous for making the future seem very much like the present. Technologies are elusive, and as new ones emerge, the efforts to get their bandwagons rolling often outstrip their realities, their deliverables. Wireless networking didn't escape the promising-technology hype and has, perhaps, suffered by the premature promises of vendors in search of the next big thing.

But all signs indicate that wireless is, indeed, a reality and that it's here to stay. Mobility has been a mantra of computing since the first portable PCs appeared, with the rebirth of paging and the proliferation of cellular phones providing a much needed impetus.

Today, it's estimated that there are wireless nets in use in businesses and homes. But the sudden spurt in growth hasn't come without the expected growing pains. Security issues, in particular, have put many companies on the slow track to mobile computing. The security shortcomings were relatively unimportant when wireless first took root in home computing environments. Essentially a case of caveat emptor, users were left to their own devices to find the holes and patch the leaks. But in the security-conscious realm of corporate computing, caveat emptor doesn't cut it.

This isn't to suggest that wireless vendors were asleep at the wheel, although perhaps they might have shared some guilt for hastily rolling out products. But, as banal as it may sound, today wireless is fully into its maturity—or at least well on its way to maturity.

New standards are addressing key issues such as security and speed. The steady stream of little letters following "802.11" seems endless, with each indicating an incremental improvement and all leading to faster, safer wireless transmissions.

The mobile mosaic

Wireless computing has, in fact, become a catchall term for a number of untethered technologies, including Bluetooth, Wi-Fi (802.11),

and digital mobile phone technology. These pieces, though disparate in function, design, and device support, roughly comprise the parts of a puzzle that could ensure wireless connectivity near and far. Bluetooth's short-range RF is the likely candidate for connecting computing devices with peripherals; 802.11-based systems offer a wider operational range that make them ideal for replacing wired access in offices and commercial venues; and mobile telephone carriers appear to have the edge for long-distance data communications.

In addition, all three have undergone significant development so that they are now poised for practicality or are already demonstrating their utility. Prices, too, have dropped to affordable levels so that Wi-Fi, for instance, can be a cost-effective or even cost-saving alternative to traditional wired networks.

There are, however, still enough unsettling aspects to keep Wi-Fi—or any of its wireless cousins—from being a slam-dunk decision for IT managers. For example, despite the efforts of dozens of vendors to address wireless security, the magnitude of the issue is underscored by a somewhat secretive meeting convening in Washington this week to address the security issue. The conclave, dubbed A Roadmap to a Safer Wireless World, includes industry representatives, academics, and government agencies such as the Department of Defense, the Department of Justice, and the National Security Agency. Ironically, news of this meeting came on the same day that retail giant Best Buy suspended the use of wireless cash registers at nearly 500 stores because a security lapse may have enabled a hacker to snag a customer's credit card number.

But both of these news items can be taken in a positive light as well. The fact that a conference on wireless security can draw such a roster of participants is a good indication that wireless computing is to be taken seriously. And even Best Buy's unfortunate foray into mobile cashiers can be spun easily to be interpreted as

an indicator that wireless networking is being taken very seriously.

It's also a sign of wireless networking's importance that it's no longer being dismissed as a fad—or worse, a gadget technology. And even more encouragement is offered by the vendors of wireless products, such as ARM's announcement of its ARM11 technology last week at the Embedded Processor Forum. New ARM processors for PDAs and other mobile devices will operate at speeds exceeding 300 MHz—the state of the art for desktop computing not too many years ago.

While there's still a lot of work to do to achieve the goal of ubiquitous, unwired computing, it is against this backdrop of both promise and peril that dozens of wireless device vendors are showing their wares at the annual NetWorld+Interop trade show and conference in Las Vegas. The sheer number of vendors and their broad product offerings are just further testament that wireless is, indeed, real and here now. ~

Look Ma, no wires!

Oct 25, 2001

By David Berlind, ZDNet

For the past six months, I have been conducting an experiment with Wi-Fi (802.11 wireless Ethernet) that has led me to four conclusions. First, no company should continue deployment of wired technology where special applications don't demand it. Second, your wireless deployment should focus on systems that have Wi-Fi built in. Third, companies must be prepared to help employees with Wi-Fi installations at home. Finally (the vendor recommendation always comes last), systems vendors must be aware that corporate IT will require Wi-Fi that is built in to the system (and not added as an afterthought through an existing expansion port). Wi-Fi notebooks from IBM and Toshiba exemplify this built-in approach. Fortunately, my experiment included neither, which allowed me to understand the pitfalls of the afterthought approach.

Wi-Fi: Go for it

I'll never forget the first LAN I had to manage. It was a 3Com 3Share Plus-based setup: an 8088-class PC server connected to a 30MB disk drive the size of a Volkswagen Rabbit. The drive interface, from a company called

Emulex, was very finicky. It had more jumpers than I could count, and the only reference manual was a yellow sticky note filled with illegibly scribbled notes and diagrams. But the worst part of that LAN wasn't its heart. It was its circulatory system: a coaxial cable whose circuitous route stretched through raised floors, dropped ceilings, treacherous precipices, and across the floor underneath every person's desk. I do not have fond memories of running around with a coaxial terminator to isolate the misbehaving segment each time the LAN went down.

When the first twisted-pair Ethernet hubs came, we didn't even wait for the standard (10Base-T) to be ratified. We just went for it. For me, the result was like breathing pure oxygen. From that point forward, the LAN went down plenty of times, but never because of a wiring problem. There were days when I would visit the wiring closet for no reason other than to pinch myself.

If twisted-pair Ethernet is your oxygen, Wi-Fi will be your nitrous oxide. There are numerous benefits to be gained from deploying Wi-Fi. For starters, simply getting someone

connected to the network is a no-brainer. No more crawling under desks, or figuring out what port in the wall to use, or going back to the wiring closet to manage contention for ports during growth periods (admittedly not a problem right now).

Just yesterday, I had to bring up a Windows XP system (in addition to the Windows 98, Red Hat, Mac OS X, and Windows 2000 boxes already in my office). It took me all of one minute to get it connected to the network via Wi-Fi. Granted, most people don't have this sort of contention for ports in their offices, but the experience was a reminder of how much easier it is to get someone on the LAN, regardless of the reason. Maybe they were just hired. Maybe they moved from another part of the building. Or maybe they're visiting from another office and need access. Wi-Fi is infinitely easier.

This ease of deployment translates into an insurance measure as well. In the wake of last month's tragic events, several stories have emerged about wireless technologies stepping up to the plate. Should you find yourself in a situation that demands the rapid deployment of an entire LAN, there ain't nothing like Wi-Fi. Not only can it be deployed faster from scratch, but if you have surviving systems that, because of your foresight, were already Wi-Fi enabled, your ETA to be up and running will be even sooner. Even if this isn't the case, just having developed the background in Wi-Fi will get you there sooner.

Don't settle for second best

My experiment included Toshiba's AccessPoint Wi-Fi solution and a bunch of notebook computers. The Toshiba solution—OEMed versions of Lucent Technologies' ORiNOCO wireless hubs and PC Cards—demonstrates why a PC Card-based solution is smooth, but not nearly as smooth as something built in to the notebook's chassis.

The wireless PC Card creates several problems. First, with a protrusion that can potentially obstruct the path of a second PC Card, it doesn't look as if it can withstand a lot of stress. Fortunately, I haven't bent or broken mine, but I've come awfully close. PC Cards with protrusions or dongles are more prone to

damage and consequently more trouble than they're worth. Second, protruding cards also obstruct the functionality of other cards with protrusions. Within days of receiving Toshiba's wireless solution, I received a test unit from the same company for a PC Card-based fingerprint reader (for biometric-based security). This card has a protrusion as well, and there is simply no way to use the two cards at the same time.

With all sorts of cards out there for all sorts of functions, Wi-Fi is something that simply has to be built in. Toshiba knows this and now offers notebook computers with built-in Wi-Fi for forward-looking corporate technologists.

For desktops, chassis- or motherboard-mounted Wi-Fi is less of an issue. Then again, I would highly discourage use of desktop systems. Notebooks cost more, but they're worth the extra expense. Not only do notebooks make telecommuting really easy, but your users will be astonished to see how their productivity goes up when they start bringing their notebook computers to meetings—without losing any connectivity. With Wi-Fi up, I leave my composition notebook behind. I bring my ThinkPad and hammer out meeting notes in much more detail than I ever could in my regular notebook or on a PDA. When a meeting's action items require a few e-mails to be sent, I usually have them sent by the time the meeting is over. Attendees without Wi-Fi will be insanely jealous when they see how much more productive you are (or very angry if you can't resist using instant messaging during the meeting).

Mixing work and pleasure

Another part of my experiment was putting Wi-Fi hubs at work and at home with the hope of moving somewhat seamlessly between the two locations. If you want to be a hero at your company, you will show the executives who already spend countless hours working from home how Wi-Fi can make it so much easier to move back and forth.

At home, I have a cable modem connected to a router, which is connected to a 10Base-T hub, which is connected to another Toshiba wireless hub. You can combine those last three

devices into one small unit using a Linksys wireless router. With Wi-Fi running in my house, moving between work and home is a breeze. When leaving the office, I simply shut the lid on my ThinkPad (putting it into suspension mode) and go home. At my house, I open the lid, and Windows 2000—smart enough to know the computer may have moved to a new network—renews its IP address. I am back on the Internet without ever connecting a wire. (Eventually, I have to plug in the power cord.)

My only complaints: I have to manually re-establish a VPN connection to regain access to the corporate network, and I have to remember to shut down Microsoft Outlook 2000 at my office and then restart Outlook in offline mode when I get home.

Once you get a taste of this convenience, you will never, ever go back. Once you give the executives in your company a taste of this convenience, you'll be a hero.


Oh yeah, security

I once read somewhere that if you had a Wi-Fi notebook, you could travel up and down Market Street in San Francisco without ever losing connectivity. The implication is that the many Wi-Fi-enabled companies along that street are keeping you constantly within range of a wireless hub.

The truth is that Wi-Fi has had some well-publicized security problems. People can access your network without your knowledge. While in a hospital waiting room the other day,

I popped the lid on my notebook to get a little work done. Much to my surprise, I was connected to the Internet. I'm sure the hospital doesn't want to be an ISP for its visitors. On the other hand, I had no malicious intent. About the only possible harm done was that I took up a bit of someone else's bandwidth.

But even with malicious intent, I'd have to be pretty sophisticated to do more harm. For starters, in every place where my wireless notebook worked, my protocol analyzer revealed that the wireless hub was behaving like a switch. (Yes, I tried, knowing that one day I would write this story.) This meant that the only traffic I could see was my traffic and broadcast traffic. I could not very easily spy on the nurse's e-mail, but if I really wanted to, I could. There are ways to sniff at wireless signals and sometimes go beyond the switch to get at other information traversing the corporate backbone. But three conditions have to exist for this to result in serious compromise. First, the person must have malicious intent. Second, the person must be pretty sophisticated. Third, there must be something worth tampering with on the backbone.

No doubt, these three conditions exist in many places. When they do, you have to think twice about deploying Wi-Fi. But there are many more situations where this simply isn't the case. Most traffic on most business networks isn't worth an outsider's time of day. But Wi-Fi is definitely worth yours. 

Use wireless technology to triumph over networking nightmares

Mar 26, 2001

By Jeff Dray

When a networking environment demands mobility or a noninvasive setup, a wireless LAN may be the best solution. While wireless LANs can't reach the speed of cable networks, a slow network is better than none at all. In this article, I will examine two cases where wireless networking was the only practical option.

Perfect for historical buildings

Wiring an older building for a network can be a nightmare. I know; I have done my share of it. In Great Britain, I'm often faced with problems that are inherent in wiring historic buildings. Running CAT5 cable through some of these buildings would be akin to cutting out the eyes of the Mona Lisa and replacing them with blinking lights. Obviously, any work involving alterations to historically listed buildings must be done with great sensitivity and care. This is exactly why wireless networks are a great way to provide older buildings with modern communications.

The museum in my old hometown is housed in a historic building upon which the Department of the Environment has placed a preservation order. Over the last eight centuries, this building has fulfilled a number of roles—none of which have made it suitable for a modern computer network.

When I was a child, this museum was dark, musty, slightly spooky, and, most of all, very

dull and boring. To liven the facility up a bit, the museum authority asked for permission to install a more up-to-date lighting system and computer network. After a long legal battle, the museum was finally allowed to install new lighting, but the authorities were immovable on the subject of data cabling. Thanks to a wireless LAN, the museum is now equipped with the right technology to move beyond being dull and gloomy.

Building a mobile network

Wireless networks are also a great option for portable applications. A local software training company that provides on-site training is a perfect example. In the past, the trainer would bring several desktops to a client's premises and network them together. This required a van and driver to carry these heavy machines around, and it also took a considerable amount of time to set up a cable network at the client's location.

Today, this company uses a wireless LAN and several laptops. The setup time has been greatly reduced, there are far fewer cables to trip over, and the whole caboodle can be loaded into the boot of a small British car. The savings in setup time has allowed the van driver to learn about computers and begin teaching, effectively doubling the company's training capacity. ~

Take advantage of the cost savings of a wireless LAN

May 15, 2002

By Del Smith, CCNA, CCA, MCSE

Without a doubt, the falling cost of wireless LAN components is a major factor driving WLAN adoption. The lower cost, coupled with a fast-maturing technology, is prompting many organizations and IT professionals to ask the question, "What is the cost of deploying a wireless LAN vs. a wired one?"

While every LAN assessment is unique, there are common factors to consider when evaluating which solution is the most cost-effective for a given situation. I'll look at both hard costs and soft costs to shed some light on how wireless stacks up against wired.

Selecting the right wireless solution

As you know, three main components make up a typical wireless LAN solution: the wireless network card, which you will find in the desktop or laptop; the access point used to connect wireless clients to the network; and the bridge, which allows for building-to-building wireless connectivity.

There are numerous vendors now offering various wireless products. You would think that vendors offering wireless network cards for around \$70 and access points for under \$200 would make the cost question a little easier to answer. But while vendors such as Linksys, D-Link, and NetGear offer inexpensive product lines of wireless products that are great for the small office/home office (SOHO) environment, you don't necessarily want to rely upon them to run a mission-critical network segment.

Businesses need to consider enterprise-class wireless manufacturers and their corresponding products. An example would be Cisco's Aironet brand of wireless products (or ORiNOCO's wireless products), which I feel are better suited to the wireless requirements of today's corporate IT environment. At the

time of this writing, Cisco's Aironet 350 series PC Card lists for about \$169, its sister PCI card for \$299, the access point is \$749, and the building-to-building bridge costs around \$1,999. (Remember: These are list prices.) At first glance, your reaction may be, "No way!" But let's take a closer look at why these products may be a better deal.

Hard costs

Of course, most of us are familiar with the costs associated with a typical wired solution. Take a couple of new corporate office buildings for example. Traditional wired costs may include CAT5 copper cable runs in the ceiling and through walls, along with their corresponding data drops needed on just about every wall feasible. I bring this up because unless you are going to run the cable yourself, quite a bit of the installation costs will be associated with laying the basic wiring and data drops.

A wireless LAN also still requires installation (preferably professional) and some degree of cabling; however, one access point can usually be installed in the amount of time it takes to terminate one data drop. To make this part of the solution complete, you may also need to throw in the cost of traditional RJ-45-based network cards, depending on whether your systems come with them preinstalled.

Don't forget about the fiber-optic cable run that may be needed to connect two buildings due to the distance limitations and conductivity of copper. Try calling your local fiber optics installer and asking the cost to connect two adjacent buildings that are 150 meters apart with fiber line. Now ask for an installation time and find out what special equipment is needed on each end. Did you mention that there's a small concrete walkway that runs between the two buildings? You'll probably be gasping for air once the installer gives you a ballpark price.

Now compare that with the costs of using two Cisco Aironet bridges to provide line-of-site connectivity between the two buildings, not to mention that these two locations can be connected and up and running just a few hours after opening the boxes. This small scenario may be overly simplified. But the fact remains that once you take into account the associated installation and setup fees, a wireless LAN can be implemented at a fraction of the cost of a wired one—and a wireless LAN can usually be set up in a much shorter time frame.

Soft costs

Remember that the most cost-effective solution does not necessarily mean it's the cheapest. There are many soft costs to consider when evaluating a wireless vs. wired network.

For starters, there's the real estate issue. If your company has a long-term lease (five or more years) or owns a building, a traditional copper-wired network could suffice for the duration of the organization's needs.

In contrast, a short one- to three-year lease may provide a greater cost value for wireless. Paying for a wired LAN in this situation could be considered a sunk cost if the organization decides to move, whereas a wireless network could be deemed an investment that moves with the company. So even when a wireless network costs more up front than a traditional wired network, that wireless network may pay for itself if you will be moving your office.

Speaking of moves, eliminating desktop Move, Add, and Change (MAC) costs is also a powerful inducement to adopt a wireless LAN. As companies downsize and upsize, they are bound to require changes in office layouts and designs. Usually, power outlets are plentiful, but data drops can be few and far between.

Another compelling benefit of wireless LAN solutions is increased mobility and productivity. Examples include doctors who can make their rounds with immediate access to patient information, conference rooms that allow access to corporate data during meetings, and libraries that enable you to complete research while remaining connected to a corporate network and/or the Internet. The list goes on and on. The increase in efficiency that can be realized by the freedom of a wireless LAN may sometimes be difficult to measure in terms of soft costs, but it's real and should be considered.

Final word

More and more organizations are leveraging their existing investment in a copper-wired network and enhancing it with a wireless LAN. This strategy offers many advantages, including the ability to add wireless “hot spots” to areas that traditionally were not wired. Many colleges and corporate organizations have implemented this in conference rooms, lobbies, and even outside working areas.

It's true that no simple equation can determine whether a wireless LAN is indeed more cost-effective than a wired one for your scenario. Both the hard and soft costs of each solution have to be evaluated, along with security, standardization, and performance issues. But with wireless prices falling and productivity gains increasing, the wireless vs. wired cost comparison deserves a closer look. Many organizations will recognize significant savings with a wireless LAN solution. ~

Get the scoop on WLANs with this wireless networking overview

Jul 31, 2002

By Brien M. Posey, MCSE

In the past two years, WLAN technology has come a long way. Prices have fallen drastically, wireless encryption protocol (WEP) security is more widely supported, and components tend to be more reliable and have a longer range; yet there are still many different factors to consider when deciding whether to go wireless. You must look at cost, reliability, speed, and of course, security.

Getting connected

Wireless networks function similarly to wired ones. However, where wired networks use cables to attach a NIC card to a hub, a wireless network uses wireless NIC cards to connect to an access point. A wireless NIC card is a NIC that's equipped with a transceiver and an antenna. An access point is a wireless hub. Generally speaking, most access points also contain an RJ-45 port that allows them to act as a gateway between a wired and a wireless network.

Technically, a wireless network doesn't require an access point. If you need only a few wireless workstations, they can run in what's known as ad hoc mode. Ad hoc mode allows a wireless NIC to communicate directly with another wireless NIC without the aid of an access point. But if you plan to use more than two or three wireless clients or if your wireless clients will require access to a wired network, you're better off running in infrastructure mode than ad hoc mode. Infrastructure mode uses an access point.

Each access point has specific capabilities that you need to be aware of. First, it has a coverage area known as a cell. Traditionally, access points have a coverage area of 150 to 300 feet in every direction. But in recent months, access points have come onto the market offering ranges of up to a mile. Special outdoor access points with large antennas can offer a range of several miles. Of course,

obstacles such as trees and buildings decrease the range and also the size of the cell. Indoors a cell's size also depends on the construction of the building. Radio signals will travel through walls, ceilings, and floors, but these obstacles can seriously degrade the signal's strength.

You must also be aware of the number of simultaneous sessions an access point can support. Just two years ago, a high-end access point typically supported about 64 sessions. Today, most access points support 256 sessions.

Multiple access points

A single access point may not be adequate for a large organization. The access point may lack the necessary range or may not support enough users. Fortunately, you can use multiple access points to add extra range and support. When multiple access points are used, the cells tend to overlap. This allows wireless users to roam from one cell to another without losing connectivity. A wireless network consisting of multiple cells works like a cellular telephone network: when a user's signal begins to fade, another access point with a stronger signal takes over.

Multiple access points can also be used for load balancing. By using multiple access points, you can split the network traffic into two or more cells, rather than having a single cell congested with all of the traffic.

Staying secure

Perhaps the biggest concern about wireless networks is security. After all, if your company uses wireless networking, someone could sit in the parking lot with a laptop and steal packets of data out of the air. This is where WEP comes in. WEP is a shared key encryption protocol for wireless networks available in 40-, 64-, and 128-bit encryption strengths. Typically, using WEP has only a small negative impact on throughput. In tests that I've

conducted, enabling 128-bit WEP seems to reduce throughput by about 300 Kbps.

Cost

Although going wireless may cost a little bit more money up front than implementing a wireless network, the wireless network will save money if the company changes locations, because the company won't have to leave behind existing wiring and go wire a new building. Instead, the company could just pick up the access points and go.

Wireless access points are actually cheaper to implement than wired hubs. At the time that this article was written, a 24-port 3Com 10/100 hub cost just under \$400. If you wanted to connect 256 users, you'd need 11 of these hubs for a total price of about \$4,400.

In comparison, wireless access points that support 256 wireless clients cost between \$200 and \$400, depending on the features that you want. Many access points even include features such as DHCP servers, firewalls, and broadband routers. Something that you must keep in mind as you look at the price difference, though, is that although wireless access points are cheaper than wired hubs, wired hubs are much faster. For example, the wired hub that I priced can run as fast as 100 Mbps. Most 802.11b wireless access points are rated only at 11 Mbps. I review a lot of access points for an independent research firm, and in real world tests, I've never seen a wireless access point run faster than 5 Mbps.

Also, wireless NICs are priced a little higher than wired NICs. A PCI-wired NIC costs about \$20 while a wireless PCI card costs about \$100. A PCMCIA version costs about \$150.

Reliability

Although wireless networks are more reliable than ever before, there are still times when reliability is a factor. Wireless networks (of the 802.11b variety) run on the 2.4-GHz frequency, just like high-end cordless phones. My network tends to slow to a crawl every time I use one of my cordless phones. I have three other 2.4-GHz cordless phones that don't cause interference problems, though.

Wireless access points allow you to select from a number of channels, but even if you happen to find a channel that always works, your network could experience the problem that I just described if the office next door gets new cordless phones.

Another possible source of wireless network interference is background radiation. For example, for a while I had a PC with a wireless NIC in my kitchen. The PC would have problems with interference every time that I used the microwave. So make sure your office kitchen isn't too close to your wireless network. ~

A primer on Wireless Application Protocol (WAP)

Jul 3, 2002

By Harshad Oak

WAP is a standard for mobile Internet applications. Its primary objective is to provide an open standard for optimized access via a mobile device to the Internet or intranet.

When first introduced, WAP was touted as a revolutionary technology that would totally transform the world of mobile computing. But WAP and WAP-based services couldn't completely facilitate such transformation due to limitations of mobile devices and mobile networks, such as:

- ▶ Small screens
- ▶ Limited device memory
- ▶ Less-powerful CPUs
- ▶ Limited bandwidth availability
- ▶ Unreliable connections
- ▶ High latency

However, there are changes on the horizon for WAP in the form of WAP 2.0. In this article, I will give an overview of WAP and how it uses WML to display content. I will also explain how WAP 2.0 improves on 1.x's features, but why you might need to continue using the 1.x standards for now.

How WAP works

When accessing a Web site from a browser on a desktop PC, the client requests data and the server sends that data in the form of HTML over an IP network. The Web browser translates the HTML data into viewable text and graphics.

On your mobile device, WAP replaces a Web browser with a WAP browser, which can also request data from a Web site. The major difference between how you access the data via a browser on your PC and a WAP 1.x browser is that the WAP browser requires a WAP gateway. This gateway functions as an intermediary between the mobile and Internet networks. When placed between a WAP browser and a

Web server, it takes care of the necessary binary encoding of content and can also translate WML to/from HTML.

Why you should use WAP

Despite initial concerns about mobile limitations, there are many good reasons to use WAP to implement mobile Web browsing:

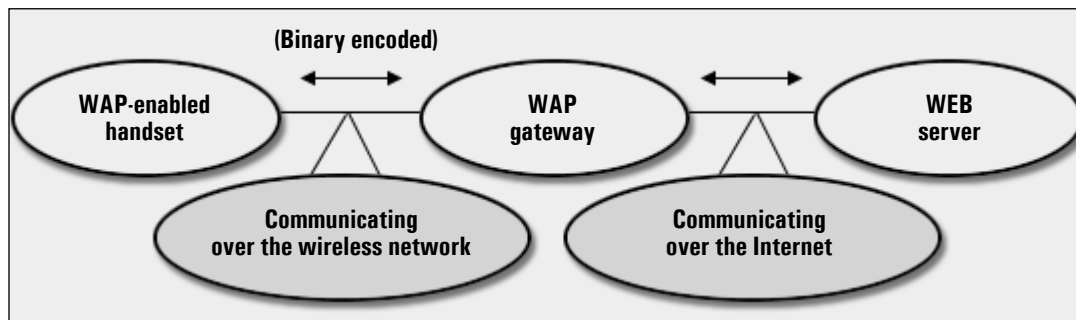
- ▶ **WAP has its own security model that works on lines very similar to Web security.** Hashing algorithms, digital certificates, and public key cryptography provide the critical security required for any real transactions using WAP.
- ▶ **WAP development is pretty simplistic.** WML and WMLScript provide for almost everything that a mobile Internet application would need. The learning curve for WML or WMLScript isn't very steep; most programmers can pick it up rather quickly.
- ▶ **WAP is widely accepted.** Major players in the wireless market (like Nokia, Motorola, and Ericsson) are all very active participants in the WAP process.
- ▶ **WAP is standard independent.** So even a switch to a GPRS network wouldn't really make a difference when browsing. Only better data transfer speeds would contribute to a better browsing experience.

WAP and WML

Wireless Markup Language (WML) is an integral part of the WAP architecture (see **Figure A**). WML is a markup language based on XML that was developed and is maintained by the WAP Forum (recently renamed the Open Mobile Alliance, or OMA).

WML is actually well-formed XML that adheres to predefined rules. It uses display tags to present content in a form suitable for mobile devices. In an ideal situation, the Web server dishes out WML content solely to be displayed on WAP browsers. A number of

Figure A



WAP architecture

WAP gateways can also translate HTML to WML. However, you shouldn't rely on this feature, because it won't really provide a truly accurate WAP display.

WAP 2.0 brings new standards

With version 2.0, WAP moved toward adopting widely accepted Internet standards. The W3C-defined XHTML Basic standard has been adopted as the basis for WAP 2.0. XHTML Basic is the mobile version of XHTML 1.0, on which the WAP Forum based its XHTML Mobile Profile.

WAP CSS is the mobile version of cascading style sheets (CSS) that has only those features of CSS that are relevant to the mobile environment. XHTML and CSS put more formatting power in the developer's command. Using XHTML and CSS, you could even dis-

play the same document on different devices using distinct presentation capabilities. WAP 2.0 also includes WML 1.x extensions to ensure backward compatibility.

With WAP 2.0, the gateway is no longer that critical a component of the WAP architecture. Also, content no longer needs to be binary encoded; XHTML goes through in text format. However, because many people still rely on mobile devices that require the WAP 1.0 standard, and because the WAP/WSP stack is being used for transport, you will still need to support WAP 1.0 gateways.

So even though WAP 2.0 offers a formidable set of features, you should probably play it safe in your current development and stick with the 1.x standards for the time being. ~

Understanding wireless LAN protocols and components

May 3, 2002

By Del Smith, CCNA, CCA, MCSE

If you listen closely, you can almost hear the sound of wireless LAN radio frequencies zipping network traffic through the air. Well, of course you can't literally hear RF waves, but wireless LANs (WLANs) are certainly being planted in IT networks from east to west. One of the most exciting technologies available today, wireless networks are being implemented by organizations of all sizes and verticals to improve productivity and decrease costs.

Understanding the different flavors of 802.11

To know where we are with WLAN solutions, we need to take a quick look at how the technology has evolved. By now, most of us have heard of the 802.11 WLAN standards established by the Institute of Electrical and Electronic Engineering (IEEE). Before 802.11, all radio-frequency wireless network communications was proprietary. 802.11 established the standards for WLANs that vendors and manufacturers follow to ensure interoperability. Entire books have been written in an attempt to clarify the various specifications and differences among the 802.11 protocol families. **Table A** briefly outlines the differences among the four.

Less confused? I didn't think so. It takes a lot more reading and research to fully understand not only the differences but also the pros and cons of each standard. The main thing to know is that the current de facto standard being adopted by most vendors and organizations is 802.11b. The next few months will more than likely reveal the slow adoption of 802.11g products based on its higher transfer rate and compatibility with existing 802.11b specifications.

WLAN components and topologies

Now, let's take a look at the typical components that make up a basic WLAN solution.

It's important to remember that wireless local area networks are just that—local. They are used within a single building or in a campus area building-to-building connection. WLANs are most often used on mobile systems as an extension to a wired LAN, as illustrated in

Figure A.

You need to be familiar with three types of WLAN components:

- ▶ Wireless network cards
- ▶ Wireless access points
- ▶ Wireless bridges

Wireless network cards come in a couple of flavors, including a PCI card for workstations and PC cards for laptops and other mobile devices. They can act in an ad hoc mode, as in client-to-client, or in a pure client-to-access-point mode. In an ad hoc mode, the wireless network card is configured to talk with other wireless network access cards that are within its range. This functionality will vary depending on the product and the 802.11 specification being used. Client-to-client (also known as peer-to-peer) WLANs are useful for small roaming workgroups of desktops or laptops that do not require access to the LAN backbone. The plug-and-play capabilities of most wireless network cards make this type of setup rather simple.

Most wireless network cards will connect to an access point. An access point is essentially a hub that gives wireless clients the ability to attach to the wired LAN backbone. The use of more than one access point in a given area is facilitated by the use of cell structures, which are similar to what cell phone providers use to maintain your coverage area.

A site survey can determine where to place access points within a building to create a map of the areas (cell structures) that will require wireless LAN access. The data transfer rate for each wireless client will be determined by its location within the cell structure. Locations

closer to the center of an access point radius will experience higher throughput than those that are closer to the outside of the cell coverage area. This is facilitated by auto shifting, which allows the data rate to downshift based on distance from the access point. Again, this functionality will vary depending on the product and 802.11 standard used.

One of the greatest benefits to roaming mobile users is the ability for one access point to hand off communication to the next access point in the roaming cell. Known as seamless roaming, this allows the user to move from cell structure to cell structure without losing connectivity to the network.

Wireless bridges enable high-speed long-range outdoor links between buildings (**Figure B**). The high-speed links between the wireless bridges deliver throughput several times faster

than T-1 lines at distances up to 25 miles. Based on line-of-sight, wireless bridges are not affected by obstacles such as freeways, railroads, and bodies of water, which typically pose a problem for copper and fiber-optic cable. Wireless bridges are often the ideal choice for campus environments where the cost of multiple T-1 lines or fiber runs can be very costly.

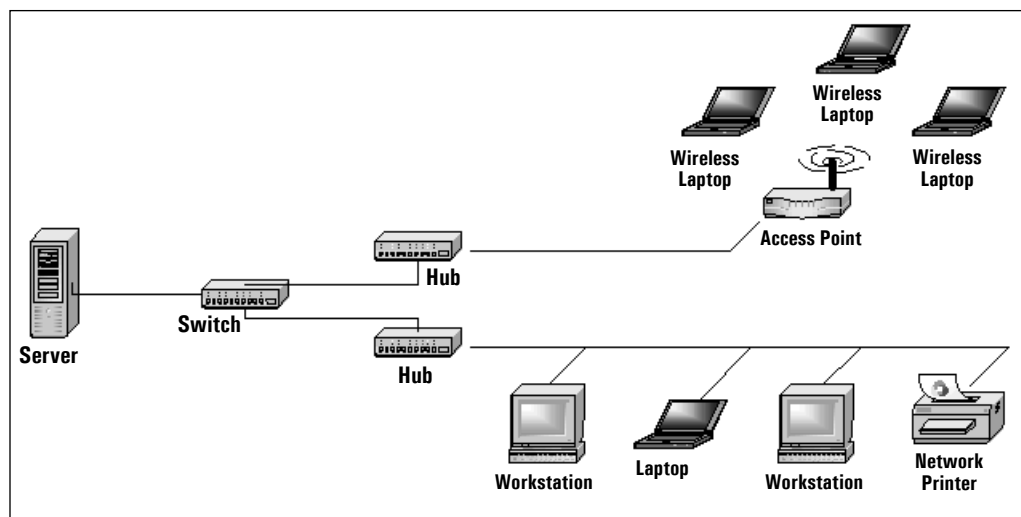
The question of wireless security

No wireless project should be implemented without a lengthy discussion of security. Over the past year, much has been written about the vulnerabilities of 802.11 wireless LANs. Older forms of security on WLANs included the SSID, which was not really a security method at all, since the SSID can easily be retrieved by sniffing the network.

Table A: Comparing WLAN specifications

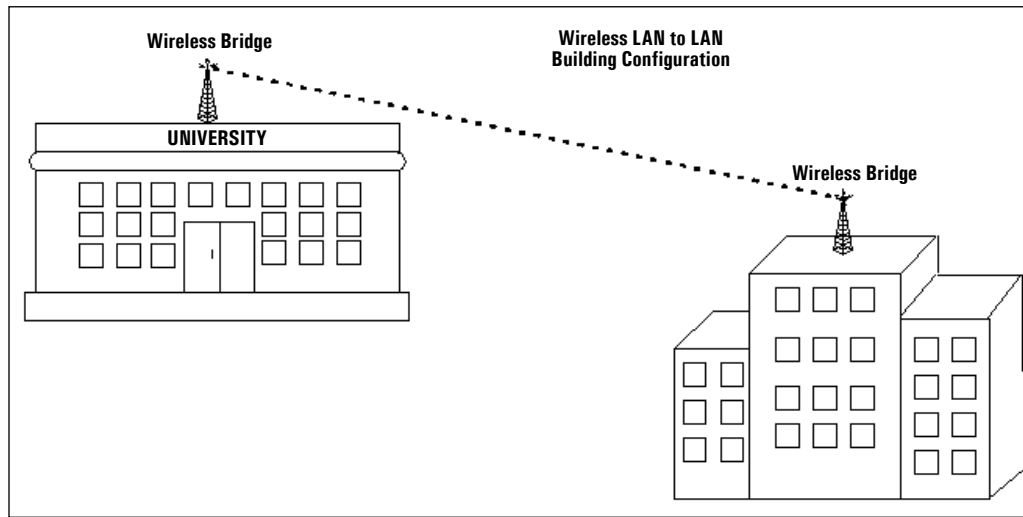
	802.11	802.11b	802.11a	802.11g
Date established	July 1997	September 1999	September 1999	January 2002—draft specification
Compatibility	802.11 only	802.11g	802.11a only	802.11b
Data transfer	1 and 2 Mbps	Up to 11 Mbps	Up to 54 Mbps	Up to 54 Mbps
Frequency	2.4 GHz	2.4 GHz	5 GHz	2.4 GHz
Modulation	FHSS and DSSS	DSSS only	OFDM	OFDM/DSSS

Figure A



This is an example of a standard wireless LAN topology.

Figure B



Wireless can also be used for building-to-building connectivity.

Authentication based on MAC filters was found inappropriate because they, too, could be sniffed on the network, and the allowable MACs could be spoofed. Newer 802.11 security uses 128-bit Wireless Encryption Privacy (WEP) for data encryption, along with shared key authentication. Unfortunately, researchers have recently identified holes in WEP that let attackers learn the keys used to encrypt 802.11b traffic.

So how does an organization protect its wireless LAN access? The IEEE has a new security standard called 802.1x that may provide the best solution. The 802.1x standard takes authentication away from access points and places it in an authentication server such as RADIUS or Kerberos. It uses the current Extensible Authentication Protocol (EAP) commonly used in PPP to control access. The

802.1x standard allows for the use of dynamically generated WEP keys on a per-session, per-user basis in place of a static WEP key placed in the access point. There are still weaknesses with this technology, and it has yet to be ratified and implemented by many vendors. So, at this time, encryption (usually in the form of VPN), traffic filtering, and other basic security restrictions on wireless network access in sensitive areas are still the best options for ensuring a secure wireless network.

Summary

As changes are in the works to establish new 802.11 standards and improve security, wireless LANs are moving into corporate America at an increasing rate. Who knows? In a few short years, wireless networks may be as commonplace as their wired counterparts. ~

Evaluating the wireless networking options

Aug 12, 2003

By Brien M. Posey, MCSE

Now that wireless networking has been around for several years and is starting to mature, companies have a variety of wireless networking standards and products to choose from. There are long-distance products used to send data between buildings miles away and then there are the shorter range products that typically provide wireless networking services within an office building or a warehouse. Both of these areas have a lot of different products and standards available, and there is no way that I could discuss them all within one article. However, since Wi-Fi is the dominant wireless networking technology at the moment, I want to discuss the various Wi-Fi options available and how to choose between them.

802.11B

802.11B is the Wi-Fi technology that has been around the longest. I implemented an 802.11B network in my home in 1999. The standard is well supported and stable. An 802.11B network theoretically supports speeds of up to 11 Mbps. However, in the real world, I have never seen an 802.11B network with a throughput above 5 Mbps. The advantages to using 802.11B are price and compatibility. 802.11B hardware is widespread and extremely inexpensive compared to 802.11G or 802.11A hardware.

There are two distinct disadvantages to using 802.11B: security and performance. Security is an issue because 802.11B is so widespread. There are numerous hacking tools designed specifically for exploiting 802.11B networks. An example of such a tool is Net-Stumbler, which detects wireless networks and uses a GPS to plot the location of each detected access point onto a map.

The biggest performance issue is radio interference. There are so many 802.11B

access points in use today that it is not at all uncommon to get interference from other access points in the area. 802.11B operates in the 2.4 GHz frequency range, which also means that it is susceptible to interference from microwave ovens and 2.4 GHz cordless phones.

802.11G

802.11G is an extension to 802.11B. Like 802.11B, 802.11G operates in the 2.4 GHz frequency range. This means that 802.11G devices are susceptible to interference from

other access points, microwave ovens, and cordless phones. So what are the advantages to using 802.11G? The primary advantage is speed. 802.11G has a maximum rated speed of 54 Mbps. To achieve the

In addition to the blazing speed, another good point of 802.11A is that it is much less prone to interference from other devices because it operates in the 5.8 GHz frequency range.

higher speeds, however, you will have to make sacrifices.

For starters, an 802.11G signal requires 30 MHz of bandwidth. The entire 802.11G frequency range consists of only 90 MHz of total bandwidth. Thus, you will be able to colocate only a maximum of three 802.11G access points within a given area.

The other disadvantage to 802.11G is range. An 802.11G signal has a shorter range than an 802.11B signal. In a way, though, this is a mixed blessing. Because of the short range, you may be able to use more than three access points to service a building, so long as no more than three access points are within range of each other at any given time.

The other advantage to 802.11G, besides speed, is compatibility. 802.11G is completely backward compatible with 802.11B. Therefore, if you already have a big 802.11B network in place and want to upgrade to something with better performance, 802.11G will allow for a

smooth transition. You would begin the transition process by swapping out the access points. Remember, though, that an 802.11G access point doesn't have the range of an 802.11B access point. Therefore, if your current access points are widely scattered or if you have wireless clients far away from the existing access points, you will probably have to install more access points than are currently in use. Once the access points have been swapped out, you can begin changing out wireless NICs. Existing clients will continue to use 802.11B until they have been given an 802.11G NIC. The access point supports both protocols.

802.11A

802.11A is a completely different animal from 802.11B and 802.11G. Like 802.11G, an 802.11A network can deliver data at up to 54 Mbps. Additionally, multiple channels can be combined for even higher data rates. I converted the wireless network in my home to 802.11A a little over a year ago. While the standard is designed for a data rate of 54 Mbps, I am using what the access point manufacturer calls turbo mode to achieve data rates of 72 Mbps. If this were a true 72 Mbps, then it would mean that my wireless network would be almost as fast as my wired network, which runs at 100 Mbps. The sad truth is that 802.11A runs more slowly than specified. While running in Turbo mode, I usually get an average throughput of about 33 Mbps on my network. Even so, that's still much faster than 802.11B.

All of this speed comes at a price. 802.11A lacks the range of 802.11B and 802.11G. The 802.11A specification provides 12 nonoverlapping channels in the 5.8 GHz frequency range. This means that you can colocate up to 12 access points. Of course, if you are using turbo

mode, you are using more than one channel, and colocation becomes more of an issue.

In addition to the blazing speed, another good point of 802.11A is that it is much less prone to interference from other devices because it operates in the 5.8 GHz frequency range. At the time that this article was written, most cordless phones operate on a frequency of 2.4 GHz. Such phones often interfere with 802.11B and 802.11G networks. 802.11B and 802.11G networks are also subject to interference from microwave ovens. At this time, not many 5.8 GHz cordless phones are in use. Therefore, because of this and the fact that 802.11A is a less popular choice than 802.11B or 802.11G, these networks are less susceptible to interference than networks operating at 2.4 GHz.

Making the decision

There are a lot of factors to consider when choosing a Wi-Fi implementation. If you are building a new network, then I recommend using 802.11A. I say this because most hackers focus on 802.11B and 802.11G networks. There are few hacking tools available for 802.11A networks because few people use 802.11A. 802.11A is also much less susceptible to radio interference than 802.11B or 802.11G because it uses the 5.8 GHz frequency range.

However, if you have an existing wireless network, you may be better off using 802.11G. 802.11G will give you the speed of 802.11A, with a much smoother transition from 802.11B. Remember that 802.11G is compatible with 802.11B. 802.11A, on the other hand, isn't compatible with either 802.11G or 802.11B. ~

Notes

Notes

Notes

Notes

Configuration and Troubleshooting

- Bridge floors and buildings with wireless access points21
- Span the WAN with wireless bridges23
- Plan effectively and save big on a wireless bridging deployment27
- Add and configure network adapters30
- Add protocols, services, and network clients and bind them all to your NIC35
- Understanding wireless network settings40
- Windows XP offers groundbreaking WLAN functionality44
- Configuring a wireless LAN connection in Windows XP46
- Create local user accounts for Windows 2K/XP peer-to-peer networking.....49
- Install a wireless connection on your home network53
- Diagnosing wireless network performance problems57
- Fix hardware and configuration issues common to wireless LANs61
- Troubleshooting the wireless woes64
- Troubleshoot wireless networking antennas66

Bridge floors and buildings with wireless access points

Aug 19, 2002

By Ron Nutter, MCSE, CNE, ASE

Lucky network administrators get to run cabling in brand new buildings. When linking buildings together, lucky network administrators with big budgets get to run fiber from building to building or maybe even get to invest in a microwave or laser connection. But not all network admins are lucky. For example, what do you do when you have to run new cabling in your 150-year-old building or link multiple buildings together on a shoestring budget? In this article, I will show you ways in which you can use wireless access points to bridge both floors and buildings together.

Since when does wireless work building-to-building?

The guidelines state that 802.11b and 802.11a communications are limited in effective range. 802.11b connections are rated only for a maximum of 300 feet, and most 802.11a connections aren't considered to be effective beyond 60 feet. Of course, the farther away you get from the access point, the slower the connection. Therefore, you may think that the usefulness of wireless communications is limited.

These limitations would imply that 802.11x connections can't handle building-to-building or floor-to-floor connections. However, as described in "Antenna on the Cheap (er, Chip)" (<http://www.oreillynet.com/cs/weblog/view/wlg/448>), 802.11b communications have been possible at distances of up to 10 miles using an antenna featuring an empty Pringles potato chip can. Longer distances (up to 20 miles) have been achieved using either commercially available antennas or something a little more substantial than a Pringles can. With this type of range, it's very easy to use 802.11b to link networks together within buildings and even in different buildings without having to run fiber optic or potentially expensive T-1 lines.

Get to the point

You have two architectural choices to consider when using 802.11b to bridge networks: point-to-point bridging or point-to-multipoint bridging. Point-to-point bridging means exactly what it sounds like; communications flow from one access point to another when connecting locations. Point-to-multipoint bridging is a little more involved. With point-to-multipoint, one central access point at your main location serves as a connecting point for all other locations/floors.

The difference between point-to-point and point-to-multipoint is analogous to the difference between Thinnet Ethernet and 10Base-T. Like Thinnet, if an access point in a point-to-point configuration fails, communication across the network will be broken, while still allowing the computers connected on either side of the break to talk to each other. In a point-to-multipoint configuration, if the failing point is one of the multiple access points, only that point's computers will lose communication. However, if the central access point fails, all communication breaks down on the network. Because the cost of 802.11b gear is very reasonable and point-to-multipoint connections create the potential for a complete network-communications failure, it doesn't make sense to try to bring all the wireless connections in on the same access point.

If you have no choice but to deploy a point-to-multipoint connection, you must take care to consider bandwidth implications. Your total incoming connections bandwidth can't exceed the maximum bandwidth of the access point. According to the 802.11b standard, the maximum top speed on a wireless network is 11 megabits per second. Assuming that you have the remote connections locked down to a 1-Mbps connection speed, this means your bandwidth would allow, at most, only 11 simultaneous connections to one access point. At this rate, each connection gets an equal slice

of the bandwidth pie available from the central access point.

If you want to use more than 11 connections, you would have to enable Quality of Service (QoS) on your central access point to ensure that each connection would achieve a specific connection speed. Because not all access points have the QoS feature, you must check to see if the feature exists on the access point you're considering before purchasing it.

If you don't want to lock down the speed of connections between access points, you can ensure that the central access point doesn't become flooded by allowing it to negotiate the connection speeds on its own. The problem with using this method is that the closer connections will obtain the faster speeds. Because closer access points would obtain higher speeds, this could cause congestion when other connection points have a demand for a higher bandwidth.

Access point placement: Floor-to-floor

Wireless access points work best for spanning between floors when the floors in question are adjacent (e.g., floors 1 and 2). If your organization is in a building where another company has an intervening floor, using wireless access points becomes problematic due to distance problems and potentials for interference.

Look for a common area on each floor where you can place the access points. A janitor or supply closet is a great second choice for access point placement if you don't have wiring closets on each floor. The locations do not have to be "stacked" one on top of the other, but stacking the access points will make the signal stronger.

Put an access point on one floor and go to the other floor to see how well you can "hear" the access point using a laptop with a wireless card. If you get anything other than a full-strength signal, move the access point to another location. If you can't accomplish a full-strength signal in any location, you will need to look at some type of directional antenna to boost the signal so that you can punch through the floor and achieve a reliable signal. Once you've found the best signal loca-

tion, place your second access point at that location. Continue on until you have access points placed on all floors of your building.

Don't forget to check what 802.11b channels are in use on both floors. Make sure the channel you use for *between* the floors isn't too close to the channels in use *on* the floors. If you have channel 2 in use on Floor 1 and channel 5 on Floor 2, using a channel between 2 and 5 wouldn't be a good idea for in-between the two floors. If you use a channel between 2 and 5, you could experience overlap and interference, which can slow down communications. In this case, you would want to use a higher channel, such as 10 or 11.

Access point placement: Building-to-building

Connecting building-to-building utilizes some of the same concepts as using access points floor-to-floor. Do a site survey at both locations with a laptop and wireless card to see what channels are in use. Once you can find a clear frequency, you will need to find a location at each building where you have line-of-sight to the target building. This may require the use of a tower to place the antenna high enough to get a clear view of the other building.


In a building-to-building connection, using the antenna that comes with the access point probably won't work. Most antennas that come with access points are typically omnidirectional, which means they send an equal signal in all directions. Because you will be sending data directly from one building to another, a directional antenna will do a much better job. Keep in mind that the higher the db gain figure for the antenna (db is a measurement that indicates how efficiently the antenna is broadcasting the signal), the narrower the beam or signal path coming from the antenna. This narrower beam means that the signal can travel farther before starting to degrade.

However, a narrow beam has one drawback: the higher the gain of the antenna, the more carefully you will need to align the antennas to get the best signal between buildings. In a building-to-building link, you should have an external (outdoor) access point. Outdoor access points are generally a little more expensive than the

access points you would use inside the building. If you don't want to use, or can't afford, an outside access point, you could also use an internal (indoor) access point in conjunction with a high-gain antenna mounted outside the building.

Regardless of type, make sure the coax cable between the access point and antenna is as short as possible. Because 802.11b uses frequencies in the 2.4-GHz range, long coax cable runs between the access point and the antenna can cause more signal degradation to occur.

Bridge the gap

Wireless networking solves a lot of problems for network administrators. Sometimes it's just too difficult to run new wires in a building, or it's too costly to connect remote buildings together on a campus. In these instances, you can use the flexibility of 802.11b to save you both time and money. Carefully place your access points, and you'll wonder why you needed cable in the first place. 

Span the WAN with wireless bridges

Jul 14, 2003

By John Kull, MCSE, Network+, A+

When faced with the task of connecting a remote office, the first options that often come to mind are a dedicated circuit (such as T1, T3, or frame relay) or site-to-site VPN. If the remote office building is near the main office, then another option is to lay fiber. However, dedicated circuits are costly and slow, and fiber is even more expensive, but yields faster connections. Site-to-site VPNs can save you money by using less-costly Internet links as the backbone for connection, but there can be QoS, security, and performance issues involved.

Before you invest in any of these technologies, you should consider another solution: A point-to-point wireless connection. Here is a look at the various methods of point-to-point wireless connectivity and how they can be implemented for WAN connections between buildings or across town.

Understanding the technology

When you think of wireless networking the first thing that probably comes to mind is the

current Wi-Fi standards: 802.11b, 802.11a, and the newest, 802.11g. The popular 11-Mbps, 802.11b standard that is typically used within a building, or building-to-building. In addition to 802.11b, several companies also use proprietary standards for wireless connections.

A typical "indoor" wireless network is made up of one or more access points that allow wireless clients to connect or associate with them. The access point provides the link, or bridge between the wired network and the wireless network. For this reason they are often referred to as a wireless bridge.

Thus, the term *wireless bridge* can be confusing. When used in the previous statement it refers to a device that connects two networks, a wired network and a wireless network. In the context of this article, we are referring to the application of connecting or "bridging" two wired networks via a wireless connection.

Unlike an access point, a wireless bridge does not connect or associate with wireless clients. It connects to another bridge device to

complete the link and join two networks together. Bridges can also be set up to provide multipoint connections enabling several remote sites (B and C) to connect to a main site (A), as shown in **Figure A**.

Bridge antennas are located outdoors, usually mounted to a roof or on a communications tower. Outdoor bridge systems use different style antennas than those required by an indoor access point. An indoor wireless network typically relies on an omnidirectional antenna, which distributes its signal in all directions creating a circular coverage pattern. **Figure B** demonstrates this concept (the black circle in the middle represents the antenna on an indoor system).

Outdoor systems use a directional antenna, such as a “yagi” or parabolic dish that focuses its signal in a specific direction, typically at another antenna (**Figure C**). An outdoor antenna typically has more power, or gain, that allows the signal to travel farther than its indoor counterpart.

Outdoor wireless considerations

Whether you are connecting two buildings 500 feet apart or five miles apart, one major consideration must be taken into account. That consideration is called line of sight. The buildings must have a location where an antenna can be mounted and “see” the other antenna.

Unlike their indoor counterparts, outdoor units do not pass signals through objects. If a tree or other physical obstruction is in the way, the signal will probably be attenuated or

reduced significantly, causing the bridge connection to fail. If an obstruction is preventing a link, multipoint hops can be installed to bypass obstacles or extend the range of a link, as shown in **Figure D**.

At a distance beyond six miles, the curve of the earth, referred to as “earth bulge,” must be taken into account (this is also seen in the diagram in **Figure D**). Earth bulge requires the antennas be mounted at higher elevations. Another consideration is the Fresnel zone, an imaginary elliptical path that surrounds the signal path. The Fresnel zone varies with distance and the frequency of the signal. The Fresnel zone must extend above any obstacles, such as trees or tops of buildings located between two points (see **Figure E**). Additional information on antenna selection and Fresnel zone considerations can be obtained from Cisco’s Web site at <http://www.cisco.com/warp/public/102/wlan/connectivity.html>.

Several other considerations must be taken into account when designing a wireless link between buildings. The speed of the link is dependent on distance. The farther away the two antennas are from each other, the slower the speed. The maximum distance is also dependent on the type of antennas selected. Cisco provides a calculation chart at <http://www.cisco.com/warp/public/102/us-calc.xls>. This is an Excel spreadsheet based on its Aironet 350 series bridge equipment that assists in selecting the proper antenna and equipment to achieve the desired speed or distance requirements.

Figure A




Figure B





Figure C



Typical distances can range from several hundred feet to 30 to 40 miles, depending on equipment selection and other factors. Cisco's calculation chart mentioned above warns that distances beyond 25 miles can pose difficulties in aligning the antennas.

Wireless bridges operate in the 900-MHz, 2.4-GHz, and 5-GHz frequency ranges. This is referred to as the unlicensed Industrial Scientific Medical (ISM) band.

LICENSING MAY BE REQUIRED

In the United States, 802.11b and 802.11g operate at 2.4 GHz, 802.11a operates at 5 GHz, and many cordless phones operate at 900 MHz. Countries outside the United States may require licensing for using wireless equipment.

For U.S. businesses, no FCC license is required to install your link. This sounds like a great thing, until you consider the fact that it means anyone can set up a link without regard to what's already installed nearby. This translates into a potential problem: interference.

An existing installation may interfere with your signal and/or vice versa. All you can do is be aware and realize you may have to relocate an antenna or change to a different channel or frequency. If you hire a contractor, ask that a site survey be completed prior to installation and prepare to resolve any interference issues that may arise with neighboring businesses.

Figure D





Figure E



Security concerns

A wireless bridge is based on the same wireless technologies as indoor wireless signals, so it shares the same security concerns. However, there are also additional factors involved that make it more difficult for wireless hackers to intercept the signal.

Since the signal is directional, a hacker would have to have line of sight to the antenna path to intercept the signal. If the antenna were located 100 feet in the air on a tower, the hacker would have to climb the tower, climb atop a nearby building, or use some form of air transportation in order to intercept the signal.

Although these considerations make it difficult for a signal to be intercepted, this is no excuse to leave the link unprotected. WEP is the bare minimum security requirement and admins should also consider additional methods to secure the signal path.

Installation

Depending on the distance involved, you may consider consulting a company specializing in

wireless bridging installations. Such a company can handle the job of erecting a tower, if necessary, and mounting and aiming the antennas.

For shorter distances, such as between two buildings on the same campus, the project may be tackled without an outside contractor. Wireless bridging “kits” are available from vendors such as 3Com and Proxim. The kits include most, if not all, components required to complete a link. Other vendors provide individual selection guides that aid in selecting the correct components based on your requirements.

Need more speed?

11 Mbps is a fast WAN connection when compared to a single 56K link, or even multiple T1 lines. However, a true 11-Mbps connection is obtainable only at short distances. But what if you want a faster connection? What if you would like to carry voice as well as data? 802.11a and 802.11g have distance limitations, which do not allow them to be used in bridge installations. In calls to both Cisco and 3Com representatives, I confirmed that neither company is offering or has plans to offer a bridge product based on 802.11a or 802.11g specifications.

Several companies do offer solutions that allow speeds above 11 Mbps and can carry both voice and data. These products are based on proprietary technologies, not the 802.11

wireless standards. Proxim offers a wide range of wireless equipment that ranges in speed from 10 Mbps to 1 Gbps. Speed is limited only by your budget. Most of its network products include a “wayside T1” for voice. Proxim’s products are based on its own proprietary standards. This makes the product less susceptible to hacking, because a hacker would have to have a matching radio to decode the transmitted signals. Additional security features are also included with its products.

Higher speed, of course, comes with a price. Proprietary high-speed wireless products are significantly more expensive than their 802.11 counterparts. However, when the cost is compared to the equivalent wired products, such as T3/DS3 service, the devices can pay for themselves over a short period of time.

Summary

Wireless links can be a cost-effective alternative to conventional “wired” services when considering a connection to a remote office or offices. Once you purchase and deploy the wireless equipment, then the only additional costs are in keeping it in good working order. Therefore, most of the cost is in the initial purchase, but it can save you money over the long run when compared to WAN services with their accompanying monthly fees. ~

Plan effectively and save big on a wireless bridging deployment

Aug 4, 2003

By John Kull, MCSE, Network+, A+

When one of our on-campus business units had outgrown its building, my boss came to the IT department and said that the company had decided to lease another building down the road and move this business unit to the new location. He asked IT to look at the options for connecting the building to the corporate network and said that they would probably need a faster connection than they have now.

The unit's current building was located on campus and connected via a wireless Ethernet bridge. The link was installed four years ago and ran at 2 Mbps. Our first thought was that a faster wireless link would work for the upgrade, but we also knew that we would have to calculate and justify the cost.

To give you a better understanding of a wireless bridging deployment, I'll explain the process we went through to estimate, prepare, and deploy this link. I'll also compare the cost savings of this solution versus a wired installation.

Planning

We first contacted a local wireless vendor and discussed various options. The vendor told us about Proxim's Tsunami wireless bridges product line, which could carry data at speeds ranging from 10 Mbps to 1 Gbps. The product line also included a feature for adding a "wayside T1" for voice.

Initially, the project was only for data, but when we learned about the voice option, we decided to extend the current voice system via the wireless link as well. The next step would be to perform a site survey to check for obstructions and to get a rough idea what heights would be involved.

For a cost comparison, we also contacted the local telco provider and discussed various options using multiple T1 lines or a partial D3 circuit. We decided initially to compare costs on a 10-Mbps leased-line connection to a 10-Mbps wireless bridging solution.

Site survey

The two buildings were located roughly a mile apart. Between the two locations lay farmland and woods. With woods come trees and trees can be a nightmare for wireless signals. To make the link work, we would have to be above the trees (since wireless bridging requires an unobstructed line-of-sight between the two units). This would require a tower at both ends of the link.

Fortunately, we already had a 150-foot communications tower located near the main campus building. After climbing our tower and visually surveying the path, the vendor recommended a 100-foot tower at the remote end of the link.

We contacted the management of the new building, who said that a 100-foot tower was out of the question. They said we could mount our equipment on the building's roof as long as it wasn't visible from the road. While doing the site survey, we went with the wireless vendor up to the roof of the new building. We located a spot in the middle of the roof, but the vendor suggested a spot near the edge that had some open area. There was just one problem: From that spot, a single tree obstructed part of the path to the campus tower. Although a tree would normally be a problem, this tree appeared to be dead. The vendor explained that the tree leaves are usually the

Figure A

Item	Telco	Wireless
Installation Costs	\$5,000	\$2,400
Equipment	\$20,000	\$22,000
Recurring Costs	\$1,600 (6 T1's = 9Mbps) \$3,600 (D3 = 45Mbps)	\$0

main transmission barrier because they contain water, so the dead tree may not cause an obstruction. The vendor representative said that they would be willing to try it but offered no guarantee.

Cost analysis

The wireless vendor submitted a quote for the installation. **Figure A** shows the costs compared to the phone vendor's quote.

As you can see, the initial equipment costs and installation were about the same for both systems. The real cost savings in a wireless system are realized because there are little to no monthly recurring charges. The only recurring cost you may encounter with a wireless installation is tower rental. In this installation, we owned the hospital tower, so we had no monthly charges there, and our initial wireless quote included the cost of purchasing a 100-foot tower for the remote location.

In addition, since the 100-foot tower was out of the picture for the remote building (because we simply mounted the equipment on the building's roof), we decided to use the money we saved there to purchase 45-Mbps equipment instead of the 10-Mbps equipment we had originally quoted for the link.

Installation

The target date for the installation was December 31, but the project was delayed due to con-

struction problems and was pushed back to February 14. This gave us more time to work but lousier weather to work in. Normally, our winters are pretty mild, but not this one. The day the antenna was installed on our tower, it was around 0 degrees. Despite the weather setbacks, the antenna and cable were mounted successfully on our main tower. The other end of the link was completed the next day.

After both ends were completed, the vendor performed some fine-tuning of the antenna alignment. The link was fired up and everything appeared to be working well. We were installed and running two weeks ahead of schedule.

Disaster strikes

My colleague and I monitored the link for the next several days. The Tsunami bridges have a Web-based management page that gives the status of the link and indicates any alarms (**Figure B**).

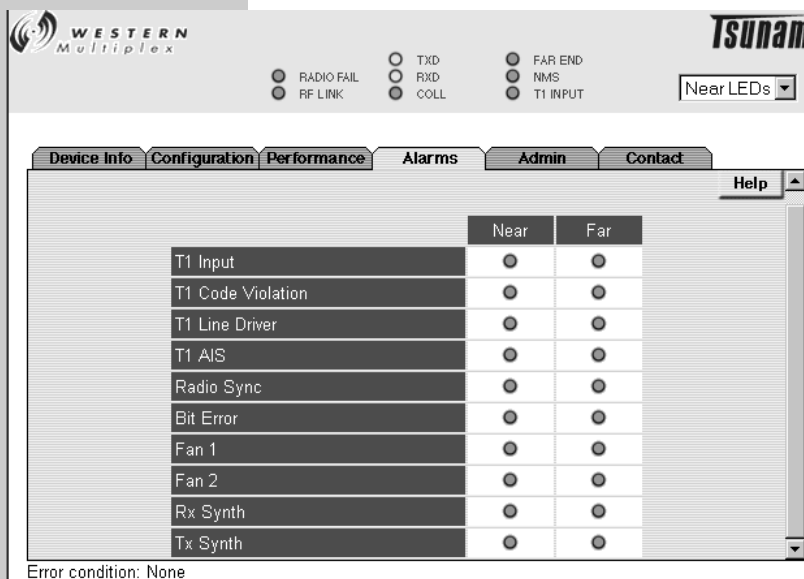
During one of those first few days, I was sitting in my office watching a heavy snow out my window and decided to check the link. Much to my surprise, it was showing a high error rate and occasionally would drop completely. I immediately placed a call to the vendor.

He explained that rain or snow would not be a problem, but if the snow was sticking to the dead tree and forming a "wall of water," that could be an issue. He promised to be out Monday to try to determine firsthand what was going on. I watched the link off and on over the weekend. It was very sporadic and didn't seem to have any correlation with the weather.

When the vendor arrived on Monday, we insisted that the antenna at the remote location be moved to the original location we liked in the middle of the roof—away from the dead tree—where we had clear line-of-sight. This location was originally ruled out because the building had a wood roof and there were concerns that it would be difficult to establish a firm and leak-proof mounting.

Nevertheless, the antenna was moved as we asked. The problem remained. We were all scratching our heads. The link had performed well for several days and then became intermittent and eventually completely dead. We

Figure B



checked connections on both ends. We replaced cables, changed lightning arrestors, and swapped the radio positions. Everything looked good. Then, we performed additional testing, and it appeared that one of the radios might be defective. One end of the link could transmit and receive to the other. But the other end could only transmit and not receive.

After battling the weather and enduring shipping delays, and after several days of troubleshooting, we ordered and installed a new radio. To everyone's surprise, the problem remained. I was starting to get some heat from the project manager. The move-in date would have to be delayed. At the direction of the manufacturer, we performed more troubleshooting and testing. The final verdict was in: We were receiving interference. A new set of radios and antenna would have to be installed using a different frequency. The current radios operated at 5.3 GHz. The new radios would use 5.8 GHz.

The new equipment was installed the next week. The weather continued to slow down the process. The new equipment worked perfectly. The link was good. We decided to allow the link to run for a least a week before we gave the okay to begin the move. We felt relieved and confident that our efforts had paid off. The data link was completed, but we had one last hurdle to clear. Now that the link was up, we could bring in the voice vendor to complete its part of the installation.

The voice vendor had no prior experience with this type of install. It had designed a solu-

tion that would extend our current PBX system to the new facility via the "wayside T1" that the wireless link provided. The Tsunami equipment is designed to accept a T1 signal at one end and make it appear at the other end. The engineers at Tsunami explained it this way: "You stick a T1 in one end and you get a T1 back out the other end." That sounded simple to us, but it was not so simple for the voice vendor. After another week of troubleshooting and tweaking, the voice system was functional. The move-in date was finally scheduled.

Final analysis

Although we had quite a few tense moments during the installation and ended up delaying the move-in date by a month, we would still choose wireless if we had to do it again. Wireless can offer significant cost savings and increased bandwidth over traditional leased-line links. As we found out, wireless requires longer lead time for testing and extra planning, and it creates more potential headaches in trying to pull off a successful installation. But, if you are willing to stick it out, the cost benefits are usually quite rewarding.

With the knowledge we gained in this installation, we went on to replace the old wireless link that had connected the original location of the business unit we moved. We performed that installation ourselves using Tsunami's QuickBridge 60 wireless kit, which included everything needed to complete a link. We completed that project in a single day. ~

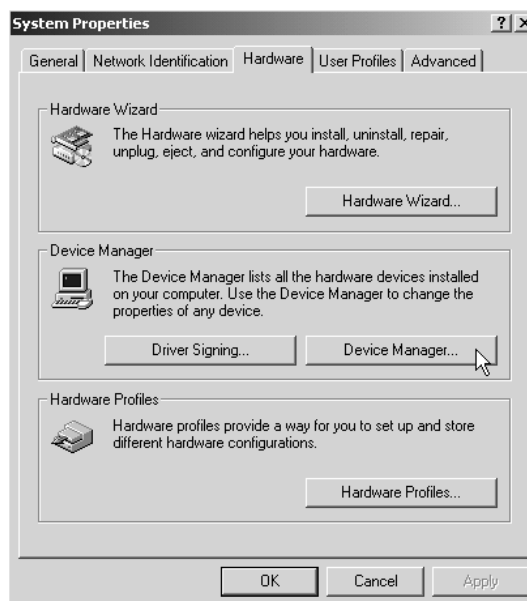
Add and configure network adapters

Oct 1, 2002

By Steven Pittsley, CNE

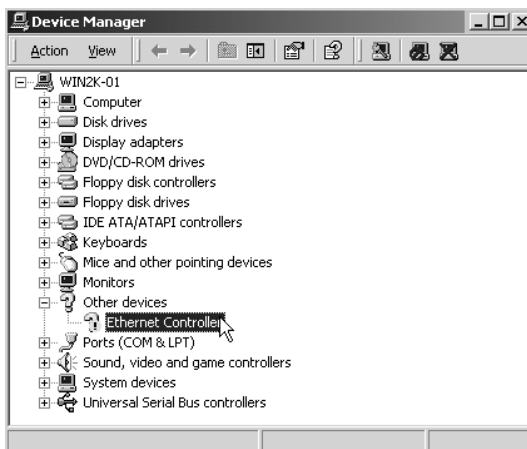
Configuring your network adapter is a fairly straightforward process that's no more complicated than installing any other peripheral device. In this article, we'll show you how to install and configure your network adapter in Windows 2000 and Windows XP.

Figure A



You'll install your NIC using the Device Manager.

Figure B



Device Manager reports a problem because our NIC has no driver installed.

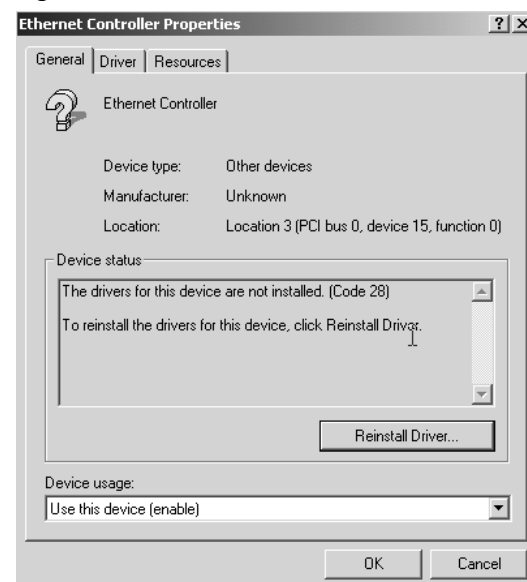
Configuring a network adapter in Windows 2000

When you add a new device to your computer, Windows's plug-and-play functionality usually recognizes the device and walks you through the installation. However, in this example, we'll show you how to manually install your network card drivers and configure your LAN connection in Windows 2000. This detailed look will help you become more familiar with the various issues involved in configuring your network connection.

After physically installing the network adapter (you'll need to reference your specific product manual for this process) you must load the software driver that allows the device to be used by your computer. To begin installing this software, go to the Start menu and open the Control Panel, then select System to open the System Properties dialog box. Click on the Hardware tab, and then click Device Manager, as shown in **Figure A**.

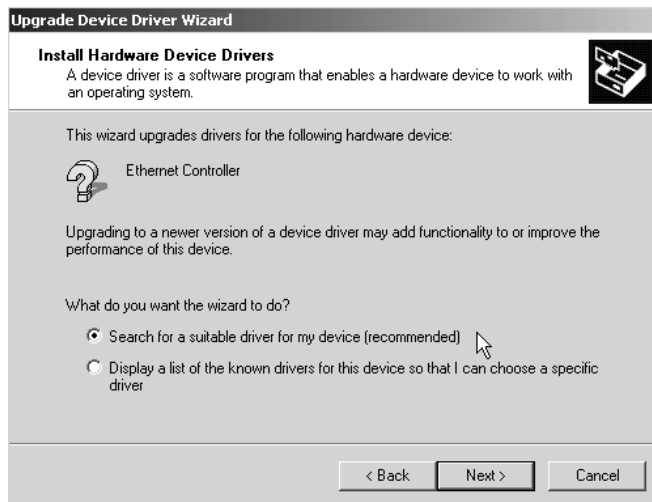
As you can see in the Device Manager window, shown in **Figure B**, the Ethernet Controller (our NIC) is listed under Other Devices. The question mark next to the controller indicates

Figure C



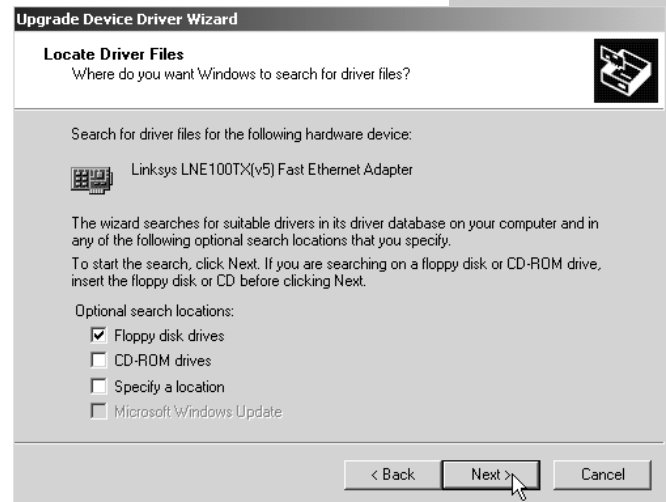
Use this interface to launch the install process.

Figure D



The wizard asks you to choose a method for finding the correct driver.

Figure E



The Locate Driver Files screen appears next.

that the device is having a configuration problem. This designation is caused because there are no drivers installed for the device.

To begin installing the software for the network adapter, double-click the device. This will display the Ethernet Controller Properties window, shown in **Figure C**.

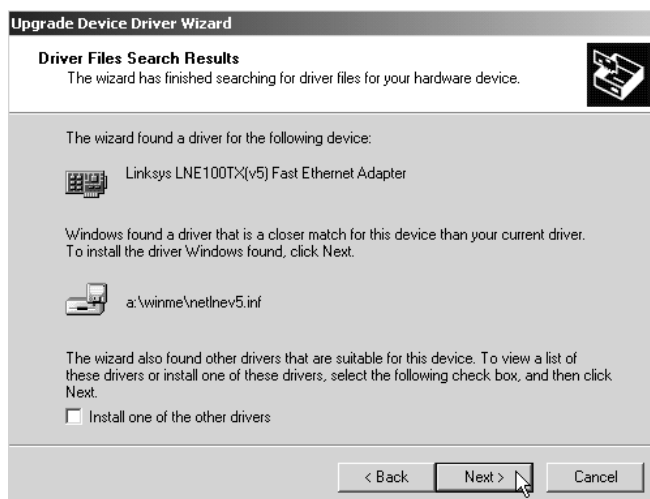
There are a couple of different ways to install the new driver, but for this example we are going to use the Reinstall Driver button that you can see in Figure C. When you click this button, you'll launch the Upgrade Device Driver Wizard. You should click Next to bypass the welcome screen and display the

Install Hardware Device Drivers screen, shown in **Figure D**.

We're going to use the default setting of searching for a suitable driver. After you click Next, the Locate Driver Files screen will appear, as shown in **Figure E**. You should select the appropriate locations where the driver files are stored and then click Next. For the purpose of this example, we're going to search for the driver files on the floppy disk drive only.

The wizard will begin its search for the driver files, and once it finds them it will present you with the Driver Files Search Results screen, shown in **Figure F**. If the correct driver is

Figure F



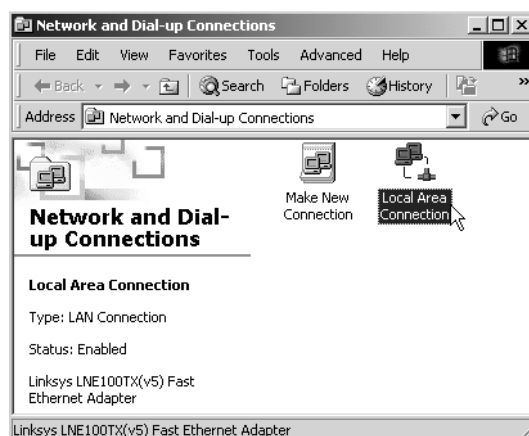
The wizard will report which driver it has found.

Figure G



Device Manager now reports that the NIC is correctly installed.

Figure H

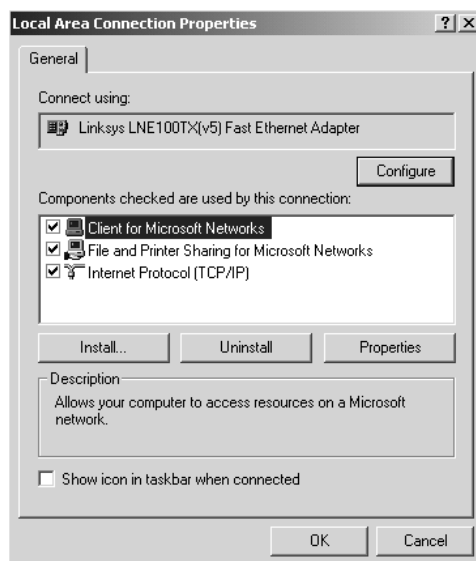


You'll use this dialog box to access your LAN settings.

found, you should click Next to continue the installation. If no drivers are found, you can click the Back button to search in a different location. If additional drivers were found, as they were in our example, you can select the Install One Of The Other Drivers option and click Next. You will then be able to select a different driver to install.

After you click Next, the driver files will be installed and the wizard completion screen will appear. You should click the Finish button to end the wizard.

Figure I



This window's options control your LAN settings.

To verify that the driver has been installed, return to the Device Manager screen. The view should automatically be refreshed, and you should see the correctly installed device under the new heading of Network Adapters, as shown in **Figure G**.

Configuring the Windows 2000 LAN settings

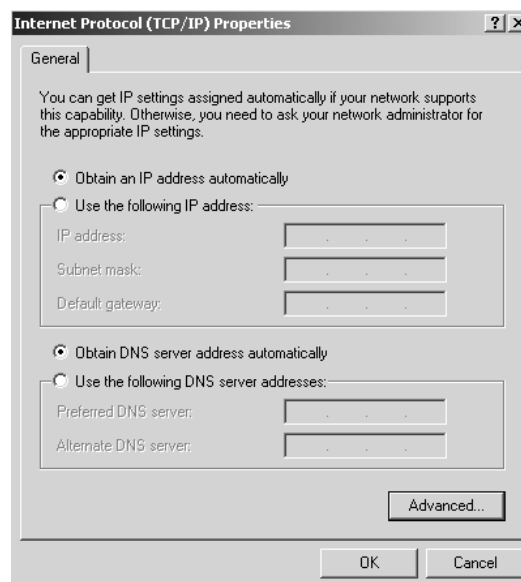
To configure your LAN settings, right-click the desktop icon My Network Places and select Properties to open the Network And Dial-up Connections window, which will look similar to the one shown in **Figure H**.

Next, right-click on the Local Area Connection icon and select Properties. This will open the Local Area Connection Properties window, shown in **Figure I**.

This window provides you with a variety of options. You can configure your network adapter using the Configure button that's located near the top of the window; you can install and uninstall networking components; and you can use the Properties button to configure the components that are already installed.

To configure your network connection, you should highlight the Internet Protocol

Figure J



Your computer is set up to look for a DHCP server by default.

(TCP/IP) networking component and click the Properties button. This will open the Internet Protocol (TCP/IP) Properties dialog box, shown in **Figure J**. The settings that are shown in Figure J are the default settings that tell the computer to look for a DHCP server to obtain a TCP/IP address (this will likely be the case if you purchase a hardware router). If you want to manually assign a TCP/IP address, you would select the Use The Following IP Address option and then enter the TCP/IP address, subnet mask, and default gateway.

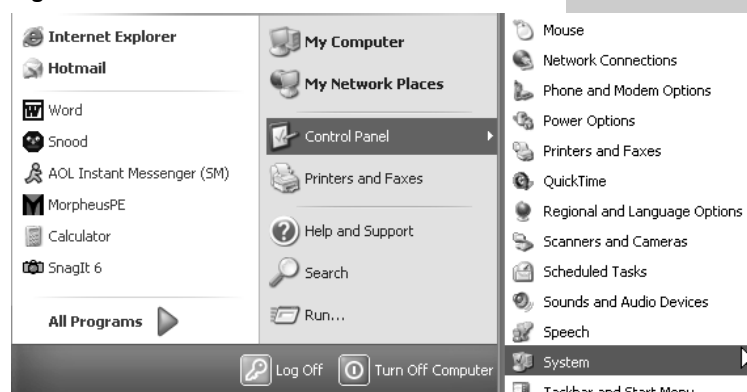
In our example, we've verified that the computer will use DHCP to obtain a TCP/IP address, and our network configuration is complete. After configuring network adapters on two or more of your network nodes, you should be able to share files and peripheral devices among the computers on your network.

Configuring a network adapter in Windows XP

Configuring your network adapter in Windows XP is basically the same process as for Windows 2000. However, the new Windows XP user interface has different ways to reach the configuration screens. In this section, we'll show you one method of reaching configuration screens so that you'll be able to configure your network adapter in Windows XP.

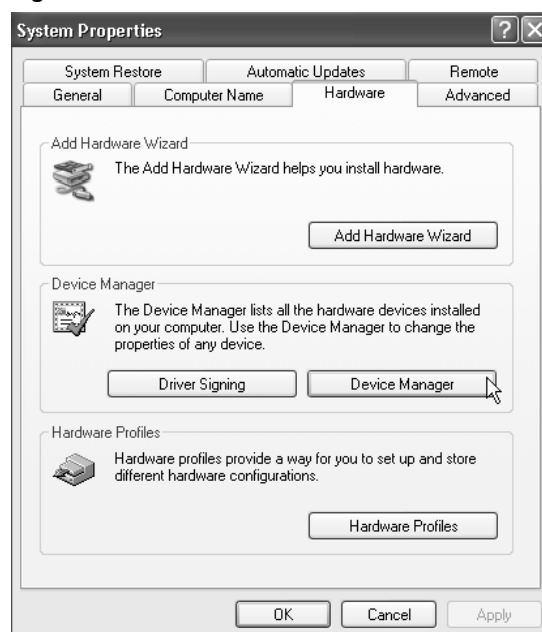
In Windows 2000, you manually installed your network adapter using the Device Manager. You can do the same thing in Windows XP, however, to open Device Manager you

Figure K



XP's new interface slightly changes the route to the Device Manager.

Figure L



XP's Device Manager is very similar to the Windows 2000 version of the tool.

Figure M



Follow this path to reach your local area connection's properties.

select Start | Control Panel | System, as illustrated in **Figure K**.

Once the System Properties window is open, click on the Hardware tab and then select Device Manager, as shown in **Figure L**. You now should be able to follow the same steps that we showed you earlier to install the network adapter driver software.

After installing the network adapter driver, it's time to configure the Local Area Connection properties and ensure that it is configured to use DHCP, in our scenario of having a router for our network. Once again, the configuration is very much the same as with Win-

dows 2000, but getting to the dialog box is slightly different. First, you must click on Start | Control Panel | Network Connections | Local Area Connection, as shown in **Figure M**.

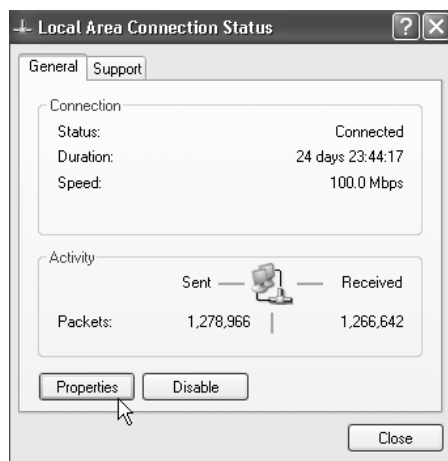
This will open the Local Area Connection Status window, shown in **Figure N**. This window provides you with some basic information about your connection, which will come in handy when troubleshooting connectivity problems.

To configure your network connection, click Properties to open the Local Area Connection Properties dialog box. You can then highlight Internet Protocol (TCP/IP) and click the Properties button. You will then see the familiar Internet Protocol (TCP/IP) Properties dialog box, where you should select the option to Obtain An IP Address Automatically to activate DHCP with your router. (If you need to manually set addresses, follow the instructions we discussed earlier in this article.)

Wrap up

As you have seen, installing your network adapter and configuring your network connection is a fairly straightforward process in Windows 2000 and Windows XP. The new XP user interface requires you to reach the configuration screens a little bit differently, but the process of configuring the devices is basically the same in both operating systems. ~

Figure N



This information may come in handy later.

Add protocols, services, and network clients and bind them all to your NIC

Oct 1, 2002

By Erik Eckel, Network+, MCP+, MCSE

The real work of creating your SOHO network begins after you've successfully installed your network interface card (NIC). To share files on a network and interconnect your systems, you'll need to make sure your system is running these three software components:

- ▶ a client
- ▶ a service
- ▶ a protocol

The process for adding these essential configuration options varies, depending on the version of Windows your system is running. In this article, we'll first run through the process in Windows 2000, and then we'll follow-up with a section on the same process in Windows XP. We'll include lots of figures to make sure you can follow along—don't worry, it's easier than it sounds.

Setting up protocols, services, and network clients in Windows 2000

If Windows 2000 was installed on your system with Typical Settings as the Networking Components option (this is common with most commercial pre-installs of the operating system), the following items should have been installed by default:

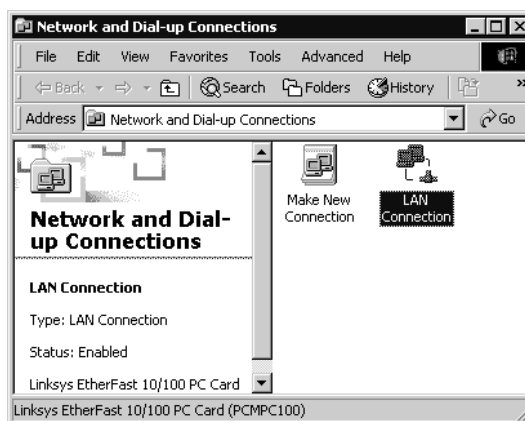
- ▶ **Client**—Client For Microsoft Networks
- ▶ **Service**—File And Print Sharing For Microsoft Networks
- ▶ **Protocol**—TCP/IP

However, instead of the Typical Settings option, the manufacturer or someone else may have specified customized settings, or someone may have had reason to delete these settings using Control Panel. If so, you'll have to reload them.

Adding the client

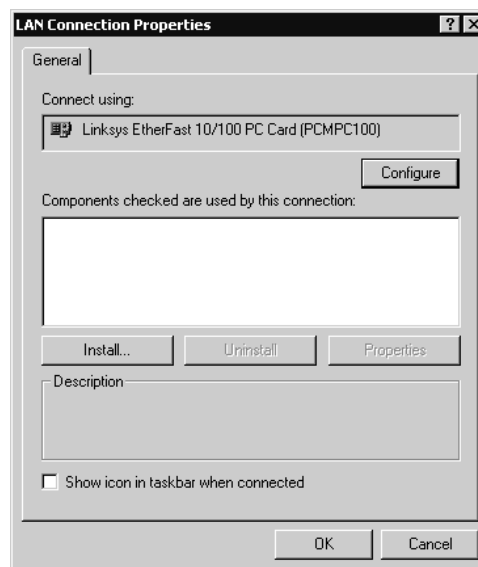
We'll start by running through the steps to install the client:

Figure A



Select the connection you want to configure and double-click it.

Figure B



The LAN Connection Properties dialog box provides information about the installed components and NIC a connection uses.

1. Click Start | Settings | Network And Dial-up Connections.
2. Select the Local Area Connection corresponding to the NIC for which you want to configure the network component settings. For this example, we selected LAN Connection, as shown in **Figure A**. Once

you've selected the connection, double-click it.

3. Click Properties in the resulting LAN Connection Status dialog box.
4. Click Install in the resulting LAN Connection Properties dialog box, shown in **Figure B**.
5. Select Client in the resulting Select Network Component Type dialog box and then click Add. Two options appear by default when using Windows 2000 Professional: Client For Microsoft Networks and Client Service For NetWare. Since we're adding a machine to a Microsoft network (that's most often the case), we selected Client For Microsoft Networks and clicked OK.
6. After the client is installed, it will appear in the LAN Connection Properties dialog box. Check the box next to it to enable its use. Click Close, and the client network component installation is complete.
7. Close the open boxes by selecting OK, then Close.

Adding the service

You install services in Windows 2000 in the same manner as you install clients. However, in

step number 5 as we described it above, select Service instead of Client from the Select Network Component Type dialog box.

You can choose from several services that are provided by default in Windows 2000:

- File And Print Sharing For Microsoft Networks
- QoS Packet Scheduler
- SAP Agent

For most SOHO networks, you'll want to select File And Print Sharing For Microsoft Networks, since that's the service that permits the sharing of files, documents, spreadsheets, and other resources, such as printers, on a Microsoft network. Then simply click OK.

Just as with the client, you'll then have to check the box to enable the service and complete this installation step.

Adding the network protocol

When you need to install or reinstall TCP/IP or install another protocol, begin by selecting Start | Settings | Network And Dial-Up Connections. Right-click the connection you want to configure and select Properties.

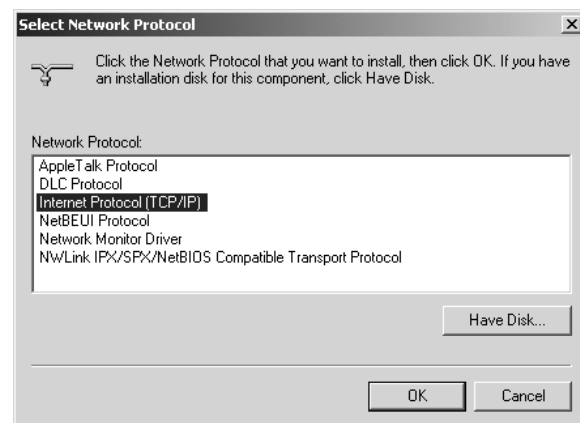
Look for the protocol you want to use in the Components Checked Are Used By This Connection box, shown in **Figure C**. If the protocol isn't listed, you'll need to add it. It's possible,

Figure C



Check before installing a protocol to ensure it hasn't already been loaded.

Figure D



In the Select Network Protocol dialog box, specify the protocol you want to install.

too, that the protocol has already been added but has been configured improperly.

Make sure the network client and service are installed (as we just did in the previous steps), then click Install. From the pop-up menu, double-click Protocol. Almost every computer, including those on your home network, relies on TCP/IP, so that's the protocol we'll install in this example.

From the Select Network Protocol dialog box, shown in **Figure D**, double-click Internet Protocol (TCP/IP) or select Internet Protocol (TCP/IP) and click OK. Both actions will select the TCP/IP option.

After installing the protocol, make sure the check boxes are selected for the network client, service, and protocol. The next step is to configure basic protocol settings.

If the system is to receive an IP address automatically from a Dynamic Host Configuration Protocol (DHCP) server, you can simply click Close and reboot. (This typically will be the case with a home network with a dial-up connection, which gets its address information from the DHCP server at your ISP. Broadband setup will vary widely, depending on your provider.) When rebooting, the system will send out a DHCP discover message. The DHCP server will snag that message off the network and fire back an IP address and subnet mask, an address for DNS services, and a default gateway (if the DHCP server is so configured).

If you want to specify a static IP address, click on TCP/IP and select Properties. In the General tab, select the Use The Following IP Address option, as shown in **Figure E**.

Once you've entered the IP address and its associated subnet mask, along with the addresses for the default gateway and DNS servers, click OK. When entering IP addresses, type periods to separate your dotted-decimal entries. Use the [Tab] key to move from box to box.

After you've provided the necessary addresses, you can specify any WINS servers you want to use by clicking the Advanced button. Click on the WINS tab, enter the WINS server address, then click OK.

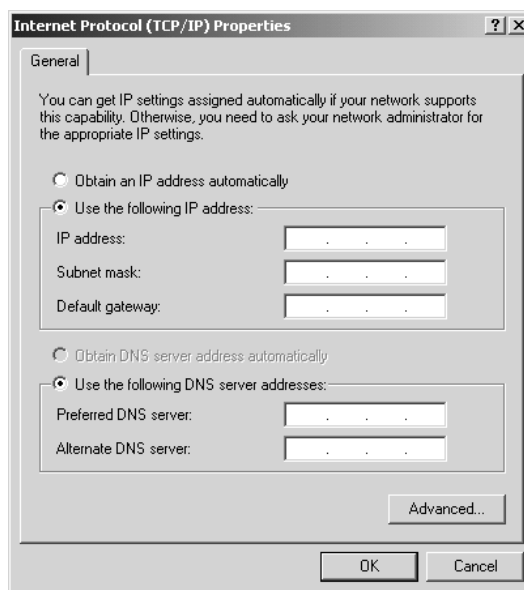
In the WINS tab, you can also specify whether to enable LMHOSTS lookup and

NetBIOS over TCP/IP. You can specify that the NetBIOS configuration be set based on a DHCP server setting. Click OK once you've set these values.

Click OK, click OK again, and click Close. Finally, close the Network And Dial-Up Connections box, and you should find your network connection working properly.

You can learn more about your network settings by running the IPCONFIG command. Do so by clicking Start | Run. Type *cmd* and click OK. Then, type *IPCONFIG /ALL*. The

Figure E



To specify an IP address, you'll have to provide the associated subnet mask, default gateway, and DNS server addresses.

OTHER AVAILABLE PROTOCOLS

Windows 2000 offers some other protocol options:

- ▶ AppleTalk
- ▶ DLC
- ▶ NetBEUI
- ▶ Network Monitor Driver
- ▶ NWLink IPX/SPX/NetBIOS Compatible Transport

These protocols typically are useful for larger networks or those that will encompass multiple operating systems, so they are usually not particularly useful in a home network setting.

details of your network adapters will be displayed. You can use this information to further troubleshoot errors on your network. It will at least lend confidence that all network adapters are configured properly.

Setting up protocols, services, and network clients in Windows XP

Just like in Windows 2000, the necessary protocols, services, and network clients should be present under a typical Windows XP installation. Again, these items are:

- ▶ **Client**—Client For Microsoft Networks
- ▶ **Service**—File And Print Sharing For Microsoft Networks
- ▶ **Protocol**—TCP/IP

Just in case you do need to add any or all of these items, here are the steps to follow.

Adding the client

Follow these steps to install the client:

1. Click Start | Settings | Control Panel.
2. Click Network And Internet Connections.
3. Click Network Connections.
4. Right-click on the Local Area Connection corresponding to the NIC for which you

want to configure the network component settings and select Properties.

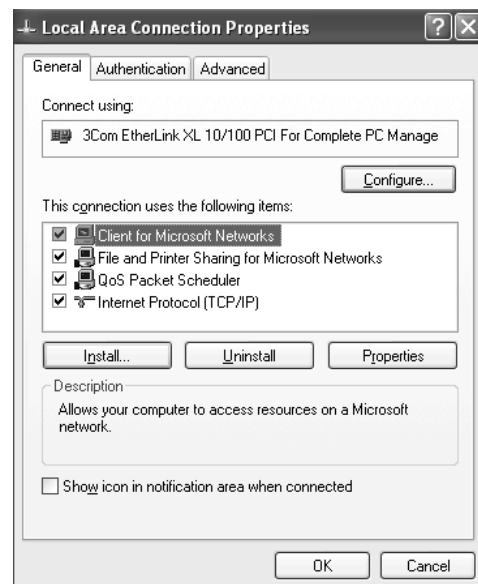
5. Click Install in the Local Area Connection Properties dialog box, shown in **Figure F**.
6. Select Client in the resulting Select Network Component Type dialog box and then click Add. Two options appear by default when using Windows XP Professional: Client For Microsoft Networks and Client Service For NetWare. Since we're adding a machine to a Microsoft network (the most common scenario), we will select Client For Microsoft Networks and click OK.
7. After the client is installed, it will appear in the Local Area Connection Properties dialog box. Select the check box next to it to enable its use. Click Close, and the client network component installation is complete.

Adding the service

You install service components in Windows XP in the same manner that you install a client component. However, select Service instead of Client from the Select Network Component Type dialog box (step 6 above).

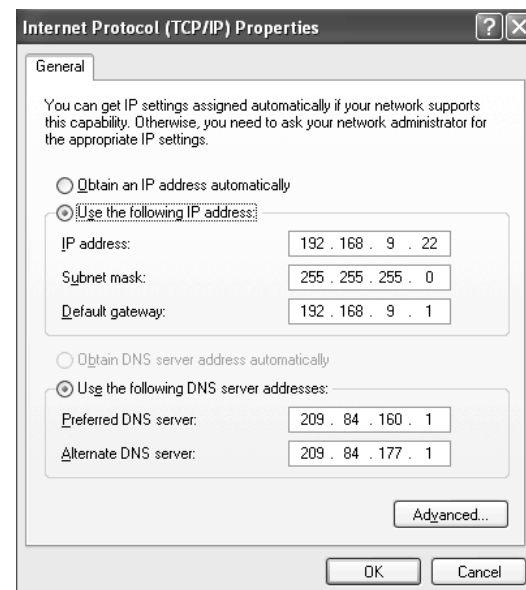
You can choose from several services that are provided by default in Windows XP:

Figure F



This dialog box displays properties for your connection.

Figure G



To specify an IP address, you'll have to provide the associated subnet mask, default gateway, and DNS server addresses.

- File And Print Sharing For Microsoft Networks
- QoS Packet Scheduler
- SAP Agent

Select File And Print Sharing For Microsoft Networks—which will let you share files, documents, spreadsheets, and other resources, such as printers, on a Windows network—and click OK. Select the check box to enable the service, and you’ve completed the service installation.

Adding the protocol

When you need to install or reinstall TCP/IP or install another protocol in Windows XP, begin by following these steps:

1. Click Start | Control Panel.
2. Click Network And Internet Connections.
3. Click Network Connections.
4. Right-click on the Local Area Connection corresponding to the NIC for which you want to configure the network protocols and select Properties.

If the protocol you want to install isn’t listed, you’ll need to add it. First, make sure that the network client and service are installed, as we discussed above. If these components are already available, click the Install button.

Next, from the pop-up menu, double-click Protocol (alternatively, you can click Protocol and click the Add button). Again, we’ll stick with the ubiquitous TCP/IP protocol for our example. From the Select Network Protocol list box, select Internet Protocol (TCP/IP) and click OK. (If you already have TCP/IP installed, this selection will not be available.)

Windows XP offers other protocols, Network Monitor Driver and NWLink IPX/SPX/NetBIOS Compatible Transport Protocol. Three legacy protocols that were available in Windows 2000—AppleTalk, DLC, and NetBEUI—are no longer available by default in Windows XP.

After installing the TCP/IP protocol, make sure that the check boxes are selected for the network client, service, and protocol. The next step is to configure basic TCP/IP protocol settings.

If the system is to receive an IP address automatically from a DHCP server (which is most likely the case with a dial-up connection but not necessarily with broadband), you can simply close the Local Area Connection Properties dialog box and reboot your system. When rebooting, the system will send out a DHCP discover message. The DHCP server will snag that message off the network and fire back an IP address and subnet mask, as well as other network settings, such as IP addresses for DNS services, a default gateway, and myriad other options.

If you want to specify a static IP address, click on Internet Protocol (TCP/IP) and select Properties. In the General tab, select the Use The Following IP Address option, as shown in **Figure G**.

Once you’ve entered the IP address and its associated subnet mask, along with the addresses for the default gateway and DNS servers, click OK. When entering IP addresses, type periods to separate your dotted-decimal entries. Use the [Tab] key to move from box to box.

After you’ve provided the necessary addresses, you can click the Advanced button and specify any WINS servers you want to use. Click on the WINS tab, enter the WINS server address, and then click OK.

In the WINS tab, you can specify whether to enable LMHOSTS lookup and NetBIOS over TCP/IP. You can also specify that the NetBIOS configuration be set based on a DHCP server setting (if you’re using DHCP). Click OK once you’ve set these values.

Click OK, click OK again, and click Close. Finally, close the Network And Dial-Up Connections box, and you should find your network connection working properly.

See, that wasn’t so bad, was it? As in Windows 2000, you can learn more about your XP network settings by running the IPCONFIG command. Just click Start | Run, type *cmd*, and click OK. Then, type *IPCONFIG /ALL* to see the details of your network adapters. You can use this information to further troubleshoot errors on your network. It can give you a quick look at network protocol settings and provide confidence that all network adapters are configured properly. ~

Understanding wireless network settings

Apr 10, 2001

By James McPherson

Wireless networking is being implemented in many IT shops. Because wireless networking is very new, however, few IT pros have had significant exposure to the unique settings it requires. In this article, I'll offer a few notes to help you set up clients and access points, discuss the settings unique to wireless devices, and detail some standard wired options that affect special features of wireless devices.

Client setup notes

Wireless network interfaces are available in PCI, USB, and PC Card formats. USB devices should be connected directly to the computer or to a powered hub because most draw their power from the USB cable. PCI and PC Card devices should be installed in a slot that provides maximum exposure to the antenna. Take care to reroute cables away from the antenna to minimize RF interference. Use shielded cables and speakers wherever possible; electrical interference will reduce your maximum bandwidth.

Client settings

When you set up your wireless clients, you'll want to carefully consider whether you should keep default settings. While these settings will get you up and running quickly, they also could compromise security. Some of these settings need to be configured on the access point as well. Make sure they're the same. Client settings include the following:

- ▶ **Ad Hoc, or Peer-To-Peer, Networking:** Some wireless devices can be set to communicate with one another without using an access point. This ability increases the flexibility of the client systems, but it can compromise a centrally administered network security policy.
- ▶ **Encryption Keys:** These keys are the values used to encrypt the data. They must match on both the client and access point. The default keys are acceptable for allowing clients to easily be added to your network,

but in a location requiring maximum security, the keys should be changed regularly to prevent intruders from breaking the encryption.

- ▶ **Mobile IP:** Cellular wireless networks allow clients to roam from one wireless access point to another. In a large enough network, this could cause a client to enter a different subnet. Normally, this would cause an IP conflict; however, the use of mobile IPs creates a kind of forwarding address, enabling access points to reroute data across subnets. Mobile IP should not be used other than in especially large continuous wireless networks.
- ▶ **Rate Control:** Rate Control allows you to specify the communication speed. Reducing the maximum bandwidth increases the roaming range and reduces power consumption but at the cost of peak performance. The defaults are usually the best general-purpose settings. This setting may be configurable to allow different default speeds in each location.
- ▶ **WEP:** The encryption scheme used by the wireless standard (802.11b) is called Wired Equivalent Protection (WEP) and is intended to compensate for the lack of physical security. Not all wireless systems provide encryption. The default for 802.11b is the internationally exportable 40-bit encryption, but some U.S. models also support the much-preferred 128-bit encryption. Sometimes, encryption is disabled by default. This option should be enabled.
- ▶ **WLAN Service Area:** This value is analogous to a network workgroup, except that clients in the same service area can communicate with one another. Configuring different WLAN service areas allows multiple wireless networks of the same type to overlap in the same geographic area. Sometimes, a service area number—for example, 101—is enabled by default. You'll want to change this setting—it is a security risk.

What if you need more than one network profile?

Because mobile devices move from network to network, your vendor's network profile utility can make or break a wireless package. This is especially true for Windows 9x laptops, which have no support for multiple network configurations. Ease of use is important, especially with small-office and home-wireless setups.

The bare minimum for any network profiler is the ability to store multiple network configurations for the same device. The most advanced profilers are also capable of changing the default printers, modem settings, area codes, long-distance codes, and shared network volumes. This functionality can provide flexibility but can also become frustrating if it's too complex.

I recommend evaluating the network profiler carefully to ensure that it meets the users' networking and usability needs. Even if the hardware is virtually bulletproof, if the end users are unhappy with the application, you'll hear about it.

Access point features

From a network-design stance, the access point is the most important component; it dictates how many clients can be served, the level of encryption, access controls, logging, network management, client administration—the whole shebang. You should choose an access point as carefully as you would your core routers.

In addition to the access point's networking capabilities, examine its physical features. The 3Com Access Point includes what it dubs a "Power Base-T" connector that enables the CAT-5 cable to provide power. This facilitates installation in locations where power is not readily available. This model includes a serial port for configuration or for operating an external modem as well. Although it seems contradictory to have a wireless network with an attached modem, this allows for custom packet routing and filtering or even setting up temporary networks (such as at a trade show where a broadband connection isn't available). Other units may include USB or even Bluetooth interfaces.

Access point settings

In this section, I'll explain the configurable settings unique to wireless access points. Other options available in access point configuration screens are reports that help you troubleshoot and tune your device.

The way you initially configure access points varies among vendors. 3Com's Access Point is configured using a crossover serial cable and a terminal client. The Proxim HomeRF wireless gateway is configured using the first wireless client that it sees. Other methods for setting up access points include Web, Ethernet, Telnet, or physical switches.

Some access points require a password. Default passwords are notoriously easy to acquire and could make your network vulnerable, especially if the device can be configured remotely. You would be wise to change the password during the configuration, but be careful when you do so. Lose it and you can no longer make changes to the device. Resetting the system to clear the password may also delete all your network settings, requiring you to reconfigure the system from scratch.

Basic settings

- ▶ **Channel:** The number of available channels will depend on the type of wireless network. In a complex setting where multiple access point zones overlap—either due to multiple service area networks or overlap of consecutive access points—you will need to ensure that the channels don't conflict.
- ▶ **Default Interface:** This is the interface used to route data where no specific rules exist. It's typically Ethernet but may be serial if you're using a modem or alternate port for routing.
- ▶ **DHCP:** This option enables DHCP services. Some models will operate only as a DHCP client, but others can act as a DHCP server. In that case, you would also need to set the valid range of IP addresses that it could assign to clients.
- ▶ **Ethernet Timeout:** The system will shut down its wireless link and disconnect clients if the Ethernet connection is disabled for a given amount of time. This feature is useful

if you have multiple access points providing redundant coverage. When the station shuts down, the clients will switch to other access points. For wireless networks with a single access point, you should disable the Ethernet Timeout option. That way, you at least retain the ability for your wireless clients to communicate with one another.

- ▶ **Interfaces:** Ethernet, PPP, and RF interfaces can be enabled or disabled. Normally, you would always leave the Ethernet and RF (wireless) interfaces active. Whether you'll use alternate interfaces will depend on your exact needs.
- ▶ **Serial Port Use:** An access point's serial port can often be configured for multiple tasks, such as packet forwarding or special routing configurations via modem or other devices. If you're using the serial port for anything other than the user interface, you'll need to provide the additional configuration data, such as the dial-out number, whether to answer incoming calls, the type of connection, and connection speeds.
- ▶ **System Password:** You use this option to change the administrator password. Do not lose the password if you change it. Resetting the device can be difficult—you may need to contact the vendor's technical support.
- ▶ **WLAN Service Area:** As with client settings, the access point WLAN service area is set by default. The entry in this field may be a network ID or workgroup name. This setting isn't very robust as a security feature because simple techniques can be used to identify it.

Advanced settings

Many of these functions are optional, either because they are of only limited use or because they provide an alternate method to the generally accepted standard.

- ▶ **Agent Ad Interval:** This setting specifies the time between requests for clients using mobile IPs. Longer times can create a lag for clients moving into a new zone.
- ▶ **Load Balancing:** Access points are often capable of load balancing where coverage

areas overlap, moving clients from a heavily loaded access point to a less active one.

- ▶ **Mobile IP:** The bandwidth overhead could be excessive when mobile IP is used to relocate to a completely different network. In that case, the home access point receives the data and then forwards it on to the new access point. This multiplies the total bandwidth needs of that client for very little gain.
- ▶ **Mobile-Home MD5 Key:** This key is used to authorize the Mobile IP identities for data rerouting.
- ▶ **Telnet/Web Server:** You can use a number of different services for access point administration. As part of your security policy, however, you can restrict the available services. If you mount the access point in a difficult-to-reach location or plan on using the serial link for other purposes, however, you will need some form of remote administration. If you can find a vendor that provides a secure remote administration method—either via a secure Web server (HTTPS) or with Secure Shell (SSH)—that would be a plus.
- ▶ **WNMP:** Enabling Wireless Network Management Protocol (WNMP) allows you to propagate changes from one access point to another and reduce management overhead.

Broadcast settings

Settings in this category control your broadcast signal. You probably won't need to alter these settings unless you're experiencing connection problems.

- ▶ **Antenna Diversity:** Diversity enables an antenna to lock onto the strongest of overlapping signals. This option is enabled by default on most systems that support it and should rarely be turned off.
- ▶ **Beacon Interval:** Access points use a timing signal to allow clients to establish connections. Locations with interference may need to adjust the beacon interval to improve connection stability, but doing so comes at the cost of performance because time spent sending beacon signals is time not spent transmitting data. The alternative

is to change the number of signals per second. Be sure to read the manual, because increasing the delay reduces the number of beacons, whereas the other system sends more beacons as the rate is increased.

- ▶ **Broadcast/Multicast Queuing:** The nature of wireless is to share channels. Sharing means taking turns, which can delay broadcast or multicast data packets. The access point can set a maximum number of delayed packets that will queue before they are given priority. Default settings are fine unless you are specifically utilizing broadcast or multicast applications. If you are, you should look more at configuring a multicast mask rather than altering the queuing.
- ▶ **Client Inactivity:** Client inactivity times provide a “grace” period for clients that have their signals interrupted. Set this rating too low and you force the client to renegotiate connections. An inappropriately long timeout, however, could tie up system resources needlessly because each access point has a finite number of clients with which it can communicate.
- ▶ **Max Retries:** This setting controls the number of times the access point will try to contact a client before it aborts the transmission.
- ▶ **Multicast Mask:** This setting allows multicast packets to bypass the queue and be given immediate delivery. It is most often used for diskless systems using network resources to boot up.
- ▶ **Rate Control:** The communication speed can be specified. The defaults are usually the best general-purpose settings. If you need to extend your coverage zone, you can reduce the communication rate. Lower communication rates can get by with weaker or lower-quality connections. By reducing everyone’s connection rate and signal strength, you also lower the odds that a nearby signal will mask a more remote one.

Security settings

In addition to encryption keys and WEP settings, you have the option of changing these access-point security settings:

- ▶ **Access Control:** This setting enables you to restrict the clients that can access your network. For instance, 3Com’s AirConnect allows you to both restrict the number of allowed clients and make specific exclusions. This is an excellent way to prevent unauthorized clients from utilizing bandwidth, though it requires extra management when adding new client systems or replacing wireless devices.
- ▶ **Client-Client Communication Zone:** Peer-to-peer networking circumvents the access point’s ability to administer a consistent security policy, but it does provide more client flexibility.
- ▶ **Encryption Administration:** For security, encryption administration can be limited to specific types of connections. Many access points support Telnet and Web administration in addition to the serial connection. 3Com can restrict encryption administration to just the serial interface, though doing so prevents you from placing the device in hard-to-reach locations, like ceilings, or from using the serial port for a modem interface.
- ▶ **Event Logging:** A variety of event logs are possible depending on the access point. The most common settings would log filtered packets, load balancing, configuration changes, Simple Network Management Protocol (SNMP) or WNMP events, and operating history. Logging should be set up to match the logging done by other network segments along with the logs unique to wireless that meet your general security model.
- ▶ **SNMP:** SNMP is an advanced feature that not all devices will support. It is not difficult to configure SNMP agents into an SNMP community, but explaining the details of SNMP is beyond the scope of this article.

To wrap up

Wireless networking includes a number of new features. This can make it harder to set up a well-configured wireless network. In this article, I’ve presented information to help you choose vendors, as well as set up and troubleshoot your client and access point connections. ~

Windows XP offers groundbreaking WLAN functionality

Apr 24, 2002

By Jason Hiner, MCSE, CCNA

Imagine that you're working on an important new project. You took your laptop home last night so that you could surf for some cool pictures to download and add to the PowerPoint presentation you created for today's meeting. This morning, you bring your laptop into work, pop it into its docking station, and make a few last-minute additions and corrections to the presentation. At 8:55, you pop your laptop out and head down to the meeting, where you hook it to the projector, make your PowerPoint presentation, and then surf through a few competitors' Web sites to give your peers a better idea of what you're talking about.

The best part of Windows XP's enhanced WLAN support is that driver and WLAN configuration are absorbed directly into XP's NIC configuration

After the meeting, you and your laptop take the half-mile walk over to the building where your CTO has her office. You meet with the CTO and give her the abridged version of the presentation, surfing a couple of competitors' Web sites to give her some examples.

Finally, at the end of the day, you take two of your company's developers out for a cup of coffee at Starbucks, where the three of you sit down—with your laptops, of course—and discuss some of the technical details of your proposal. Unfortunately, one of the developers forgot to print out an important document that the three of you were going to discuss. No problem. You simply make a VPN connection to the office and grab the document off the file server and then you e-mail it to the other two developers, who receive the file in less than a minute.

In this scenario, you roamed across four networks in five physical locations. If your laptop had been configured with Windows XP and a wireless network card, you would have had network connectivity at each stop and, better yet, you would not have had to do any reconfiguration as you roamed to each place. Of course, this assumes that each location had connectivity to a wireless access point, but with the rapidly declining prices of wireless hardware and the adoption of WLANs in corporations and public spots such as Starbucks, this is definitely a plausible scenario.

Wireless LANs in Windows XP

The kind of network roaming depicted in this example would have been much more difficult (impossible in most cases) in Windows 2000 and other versions of Windows. That's because in Win2K, wireless networking configuration is handled primarily by third-party utilities that are installed along with WLAN network card drivers that come from WLAN vendors. The best part of Windows XP's enhanced WLAN support is that driver and WLAN configuration are absorbed directly into XP's NIC configuration, and WLAN network roaming is handled with precision and simplicity.

Here are the three major improvements that make WLANs work so well in Windows XP:

► **Zero configuration**—The third-party drivers and WLAN configuration utilities used with previous versions of Windows can be described as inelegant, at best. Windows XP makes the process much simpler by automatically recognizing almost all WLAN network cards (eliminating the need for third-party drivers). To configure the WLAN, you simply go into the Properties for the network card, where you will automatically find an extra tab named Wireless Networks. There you can choose from

XP UPGRADE CAUTION

One word of warning about zero configuration: If you have a pre-Windows XP system on which you have installed a WLAN driver and utility, you need to uninstall that software before you upgrade that system to XP. Otherwise, there can be some conflicts, and you will probably encounter some errors and problems when attempting to use your WLAN card in XP.


among available networks or manually configure preferred networks. This network configuration is smart, too. For example, it automatically detects when a wireless access point changes its channel ID, and if the system plugs into a 100Base-TX landline connection, it tells the system to use that connection rather than the slower (11-Mbs) WLAN connection.

- **WLAN roaming**—Our scenario showed an example of the kind of roaming that's possible with the combination of WLANs and Windows XP. Multiple preferred networks can be configured in the XP Wireless Networks tab. This can even include options in which some of these networks use static IP addresses, while others rely on DHCP. Of course, the real coup is the fact that you do not have to reboot your machine, select any menu options, or perform any configuration activities. Once you have WLANs specified in your preferred

networks, you can leave your laptop running and simply move from one WLAN network to the next. Your laptop will automatically change network configuration.

- **Better and easier security**—Of course, no conversation about WLANs is complete without giving some attention to security. Fortunately, Windows XP also builds in measures that can make WLANs more secure and that greatly simplify security configuration for administrators. XP implements support for Wireless Encryption Privacy (WEP) and IEEE 802.1x, which provides port-based, authenticated network access for wireless networks (although it can also be used for standard wired networks). Basically, the latter is built in to network card configuration, and it makes it easy to configure RADIUS authentication, smart card authentication, certification management, and other standard security protocols that handle identity management and keep intruders from being able to infringe on corporate WLANs.

Bottom line

Windows XP takes WLANs to the next level of functionality in a way that no single WLAN vendor ever has. Better yet, XP does not care what brand of WLAN network card you are using. It recognizes virtually every WLAN card available and simplifies their configuration into standard operating system menus. 

Configuring a wireless LAN connection in Windows XP

May 1, 2002

By Jason Hiner, MCSE, CCNA

I love it when things work like they're supposed to!" That has long been my favorite little catch phrase when setting up and configuring new IT solutions. Sadly enough, that phrase has become even more special to me because it's so rarely that I actually get to say it when working with today's technologies.

However, I was able to enthusiastically utter this phrase when configuring a wireless LAN connection using Windows XP. As I wrote in "Windows XP offers groundbreaking WLAN functionality" (page 44), the most valuable new feature of Windows XP is the way that it seamlessly handles WLAN configuration and roaming. Now it's time to walk you through the process of setting up a WLAN network card in XP to prove just how intuitive it is.

Install the WLAN network card

Of course, the first thing to do is pop a WLAN network adapter into your system—and it's still best to do this while the system is shut down. In most cases, you'll probably be putting a PC Card adapter into a laptop system. However, there are also PCI and USB adapters for desktop systems.

For this example, I am installing an ORiNOCO Gold PC Card into a Dell laptop. I chose the ORiNOCO card because it had good reviews from industry experts and buyers, and I was happy with the choice; the card proved to have excellent range while holding a strong signal. I highly recommend the card for corporate installs.

In my case, Windows XP was already installed on the system before I added the WLAN network adapter, but for the purposes of this tutorial, you will achieve the same effect by installing the WLAN card before loading Windows XP. If you had already installed a WLAN card (and its drivers and utilities) in a previous version of Windows, and you are now

upgrading to XP, you need to watch out for a gotcha: Before upgrading to XP, uninstall the drivers and utilities that came with the WLAN card. If you don't, you could run into some errors and conflicts with your WLAN configuration when you upgrade to XP.

Verify that XP recognizes the WLAN card

Once you power on your system, Windows XP should automatically recognize your WLAN card. (It has a vast database of WLAN adapter drivers built in.) After the card is recognized, Windows will automatically add it to the list of available interfaces in Network Connections. To verify this:

1. Click Start | Control Panel.
2. Click Network And Internet Connections.
3. Click Network Connections.

You should then see an icon that says Wireless Network Connection. Double-click that icon to bring up the Wireless Network Connection Status dialog box (**Figure A**). This should look familiar. It's basically the same as the Local Area Connection Status dialog box you see when you double-click on a standard Ethernet NIC, but there's one distinction. The wireless version has a nice little graphic with green bars to show the signal strength of your radio wave connection.

Configuring wireless networks

When you're ready to configure your WLAN settings, click the Properties button. This will bring up the network settings properties (**Figure B**) that you're probably familiar with. They're the same as the network properties for a standard Ethernet NIC, but with one important addition: When you are configuring a WLAN network card, you will see a tab called Wireless Networks.

Click on this tab, as we've done in **Figure C**. Now you can configure your WLAN adapter

to connect to various wireless access points (WAPs).

First, you'll need to make sure the Use Windows To Configure My Wireless Network Settings check box is selected. (This is the default setting.) You'll notice that there are two sections to this tab: Available Networks and Preferred Networks. In the Preferred Networks section, you can manually set up a connection to a WAP by clicking the Add button. You can then enter the Network Name (SSID) for the access point and set up Wireless Encryption Privacy (WEP), as shown in **Figure D**.


Another way to connect to a WAP is to click the Refresh button in the Available Networks section. Windows will go out and look for nearby access points and give you a list of them. Just click on the one you want to use and then click Configure. This will pull up the same Wireless Network Properties screen that you saw in Figure D, only the network name will automatically be displayed. After you tinker with the settings and click OK, the WAP will be placed on your list of Preferred Networks.

Now when you roam to new locations (especially ones that you'll probably be returning to later), you can simply let Available Networks find the access points, and you can add them to your preferred networks with a few clicks. When you return to that location, your laptop should then automatically connect you to the WAP, and you'll have network access without having to do any special reconfiguration.

If you have multiple access points in a single location, you can add them all to your Preferred Networks list and simply use the Move Up and Move Down buttons to prioritize them.

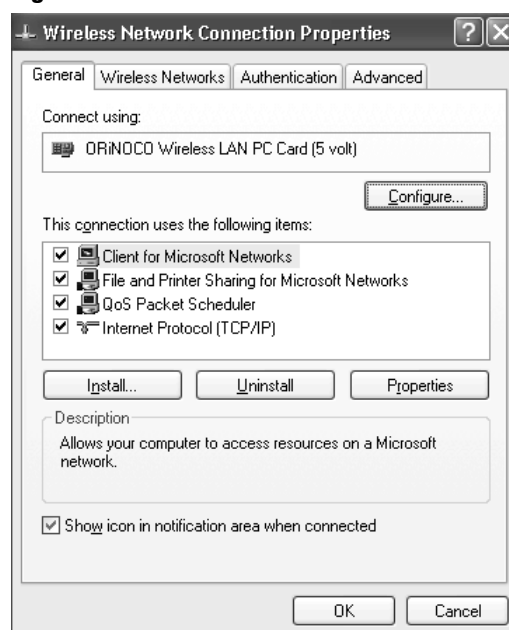
There's one more setting you should be aware of on this screen, which you can access by clicking the Advanced button. Here, you set your preference in terms of connecting to WLANs powered by access points or connecting to peer-to-peer WLANs (basically just connecting to other client machines that have WLAN network adapters installed). You also have a third option of connecting to Any Available Network, which will show you both of these categories. Obviously, in a corporate

Figure A



The WLAN status box shows the signal strength of the wireless connection.

Figure B



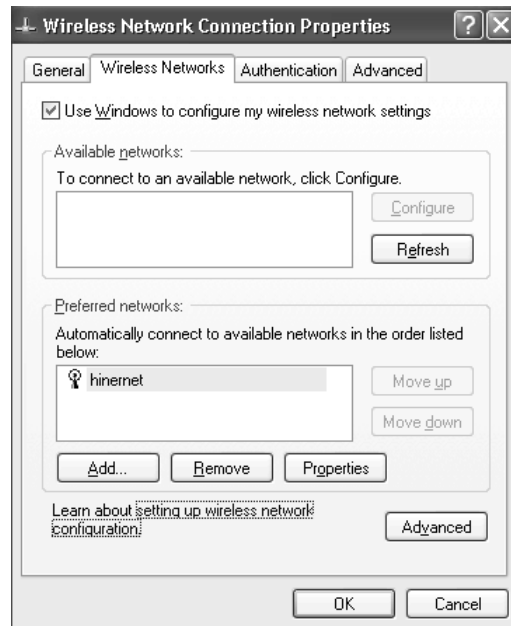
WLAN adapters have an additional configuration tab, Wireless Networks.

environment, you'll probably want to rely on access points. You'll also probably want to leave the Automatically Connect To Non-Preferred Networks check box deselected.

WLAN authentication and security

Another nice feature of the Windows XP implementation of WLANs is that it has

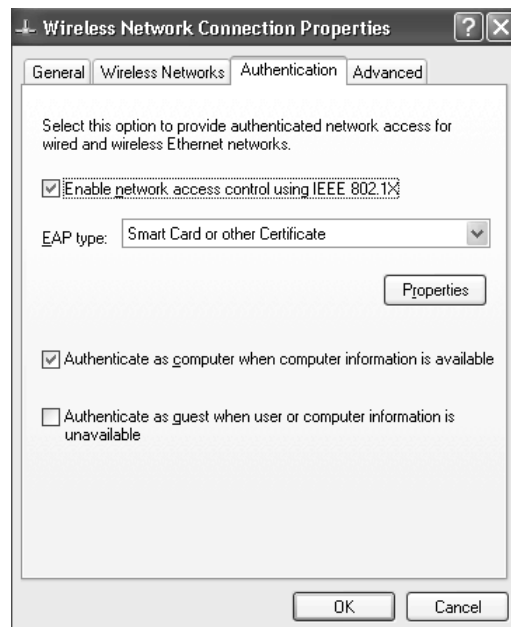
Figure C



The Wireless Networks tab is where you handle WLAN setup.

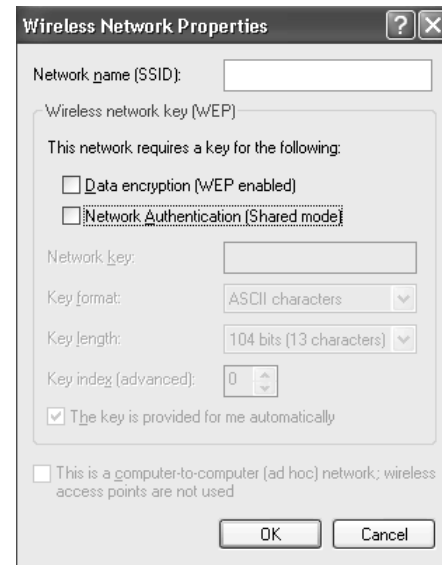
built-in support for IEEE 802.1x security. This makes it easy to require identity verification for WLAN adapters via a variety of standard authentication mechanisms, including RADIUS, smart cards, and certificates. This can be configured on the Authentication

Figure E



The Authentication tab makes it easy to configure 802.1x security.

Figure D



The Wireless Network Properties screen enables you to set up a connection to an access point.

cation tab (Figure E) of the network adapter's properties page.

It's important to note that 802.1x security is not limited to WLANs. It can be used for standard 10/100 Ethernet connections as well.


Basic monitoring and troubleshooting

Once you make your WLAN connection, you can easily monitor the reception and bandwidth of your connection. First, go into the properties of your WLAN network adapter (which appears in Figure B). Then, select the Show Icon In Notification Area When Connected check box. This will put a small icon with two computers in the system tray (in the lower-right corner of your screen). The icon will change colors when data is being sent over this network interface. (The little computer screens change from navy blue to sky blue when data is moving.) When you hover your mouse over this icon, you'll see a screen tip displaying connection information. This includes the name of the wireless network that you are connected to (usually the WAP), the connection speed (in Mbps), and the signal strength of your radio wave connection (from Very Low to Excellent).

Summary

All in all, Windows XP greatly streamlines the configuration and implementation of WLANs. In addition, it improves functionality (especially roaming) and makes it easier to implement security features such as WEP and RADIUS. To my surprise, I even found that the WLAN client software that's built in to XP is superior to the third-party drivers and utilities that come with WLAN cards for use in older versions of Windows. I found that in XP, the WLAN cards have an easier time locating and holding wireless connections, and they

don't suffer from as many inconsistencies and hiccups.

I have not been a huge fan of XP. However, its WLAN implementation is the one area where XP is head-and-shoulders above all previous versions of Windows client operating systems. If you want to configure laptops for extensive use of WLANs, you should definitely consider upgrading them to XP, especially if they are going to be roaming among different access points and/or different physical locations. 

Create local user accounts for Windows 2K/XP peer-to-peer networking

Apr 16, 2003

By Greg Shultz

Sharing resources on Windows 9x/Me systems is as easy as opening Network Neighborhood and double-clicking the share name, and maybe typing a password if the resource is password protected. However, the process is a bit more complicated when setting up a peer-to-peer network that includes Windows 2000 or Windows XP systems because of the newer operating systems' increased focus on security.

To enable W2K Professional peer-to-peer networking, you'll need to manually add additional user accounts. Fortunately, the procedure is a little easier in Windows XP because of its Network Setup Wizard. Let's take a closer look at these two approaches.

Working with Windows 2000

If you're creating a peer-to-peer network composed of only Windows 2000 systems or a mixture of Windows 2000 and Windows 9x/Me machines, you'll need to manually add additional user accounts to the Windows 2000

systems. W2K Professional was designed to work in a domain-model network where all users are verified by a domain controller. When you set up W2K Professional systems on a peer-to-peer network, there's no domain controller, of course, but users still must be verified before they can access shared resources. So, you need to create local user accounts on your W2K Professional system for every computer that will need to access shared resources on that system.

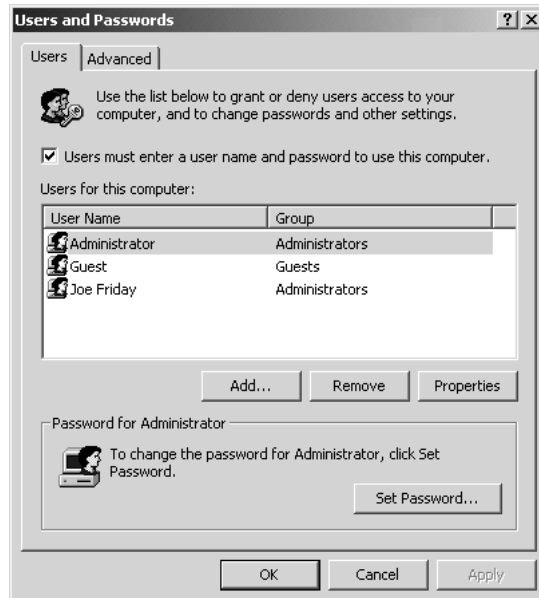
Before you get started, you'll need to create a list of the user account names and passwords on all systems on the peer-to-peer network. Once you have the list, you're ready to set up your accounts.

To begin, open Control Panel and double-click the Users And Passwords icon. When you see the Users And Passwords dialog box, shown in **Figure A**, click the Add button.

From this point, simply follow the directions in the Add New User Wizard to create an

account with one of the usernames and passwords on your list. When you get to the last page in the Add New User Wizard, you'll need to specify the level of access for the new user account, as shown in **Figure B**. The access level

Figure A



You need to set up user accounts on the Windows 2000 system for every user on the peer-to-peer network.

you choose will depend on how much control you want the user to have; in most cases, a Standard user account will be sufficient.

When you click Finish, you'll return to the Users And Passwords dialog box, where you'll see the user account in the list, as shown in **Figure C**. Repeat these steps to set up the other user accounts. When you have finished, all users will be able to seamlessly connect to the Win2K Professional system and access shared resources.

Working with Windows XP

Microsoft realized that the demand for peer-to-peer networks is on the rise, so it made creating such network configurations as easy as possible with the Windows XP Network Setup Wizard. Basically, you launch the Network Setup Wizard on a Windows XP system and follow the onscreen instructions to configure a Windows XP system for peer-to-peer networking. When you get to the last step, you have the option to create a Network Setup Disk, which you can then use to configure Windows 9x/Me systems to participate along with Windows XP. Let's take a closer look at the procedure.

Figure B



When you get to Add New User Wizard's last page, specify the access level for the user account.

Figure C



The new Sam Saturday user account allows this Windows 98 user to seamlessly access shared resources on the Win2K system.

USING MICROSOFT'S SOHO NETWORKING CHECKLIST

Before you run the Network Setup Wizard, you might want to investigate Microsoft's "Steps for Creating a Home or Small Office Network" checklist (http://www.microsoft.com/windowsxp/home/using/productdoc/en/default.asp?url=/windowsxp/home/using/productdoc/en/hnw_checklistP.asp). To find it and other helpful networking information, check out Microsoft's "Windows XP Networking and the Web" page (http://www.microsoft.com/windowsxp/home/using/productdoc/en/default.asp?url=/windowsxp/home/using/productdoc/en/hs_networking_web.asp).

On your Windows XP system, open Control Panel and select the Network And Internet Connections category; then click the Network Connections icon. When you see the Network Connections window, select the Set Up A Home Or Small Office Network item on the Network Tasks Explorer Bar to launch the Network Setup Wizard.

The first two pages of the wizard contain helpful information that you should peruse. The page you'll see next depends on whether your peer-to-peer network already has an existing shared Internet connection. If it does, you'll see a page that prompts you to use the existing shared Internet connection. If it doesn't, you'll see a page asking you to choose an Internet connection method or to configure a network without an Internet link.

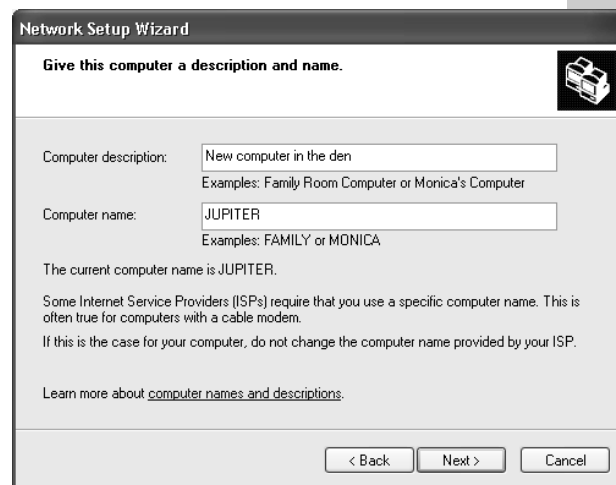
Once you work through your Internet connection options, you'll see the Give This Computer A Description And Name page. At this point, you'll assign a computer name to your system, as shown in **Figure D**.

When you click Next, you'll be prompted to specify a workgroup name, as shown in **Figure E**. If you have an existing workgroup name, just type that same name here.

When you click Next, you'll see a summary screen that shows you the selections you've made so far. When you click Next, the wizard will apply your settings and configure your Windows XP system to participate in a peer-to-peer network. Once the configuration operation is complete, you'll see the You're Almost Done page, as shown in **Figure F**, and be prompted to create a Network Setup Disk.

Even though you may not need to use a Network Setup Disk, I suggest you go ahead and create one anyway, just to have it on hand.

Figure D



You'll need to assign the system a computer name.

Figure E



You must specify a workgroup name for the peer-to-peer network.

At this point, your Windows XP system should be able to see and access shared resources on all other computers on the peer-to-peer network via My Network Places.

Likewise, all the Windows 9x/Me systems on the peer-to-peer network should be able to see and access shared resources on the Windows XP system.

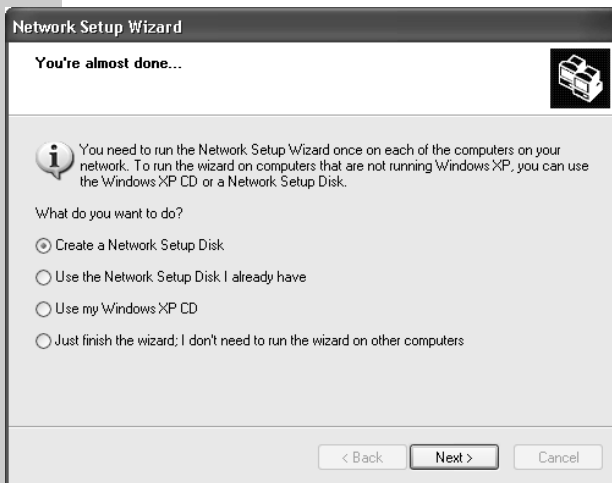
If that's not the case, you'll need to use the executable file on your Network Setup Disk to run the Network Setup Wizard and configure your Windows 9x/Me systems to participate in the peer-to-peer network. Remember that the Network Setup Wizard can run only on Windows 9x/Me systems. If your peer-to-peer network contains Windows 2000 systems, you'll need to follow the steps we covered earlier.

More information

If you want to learn more about peer-to-peer networking with Windows 2000 and Windows XP, you should investigate the following Microsoft Knowledge Base articles and Web pages:

- ▶ “Configuring Windows 2000 Professional to Work in a Peer-to-Peer Workgroup” (<http://support.microsoft.com/default.aspx?scid=kb;en-us;258717>)
- ▶ “Using Crossover Cables in Home or Peer-to-Peer Networks” (<http://support.microsoft.com/default.aspx?scid=kb;en-us;278870>)
- ▶ “Troubleshooting Home Networking in Windows XP” (<http://support.microsoft.com/default.aspx?scid=kb;en-us;308007>)
- ▶ “Contents and Function of the Home Networking Wizard Setup Disk” (<http://support.microsoft.com/default.aspx?scid=kb;EN-US;262148>)
- ▶ “Share All Your Home Computing Resources” (<http://www.microsoft.com/windowsxp/home/evaluation/overviews/connectedhome.asp>)
- ▶ “Windows XP Networking Features and Enhancements” (<http://www.microsoft.com/windowsxp/pro/techinfo/planning/networking/overview/default.asp>)

Figure F



As the last step of the process, you should create a Network Setup Disk.

Install a wireless connection on your home network

Oct 16, 2002

By Greg Shultz

Installing a wireless connection on your home network is a lot easier than you might think. Of course, the amount of work involved depends on whether you're adding the wireless connection to an existing network or building a home network from scratch that will include a wireless connection.

However, once you get down to the basics of your wireless connection, the settings are very similar, if not identical. You may need to change only a few configuration settings to get the wireless portion of the connection to work. In most cases, your wireless device's default settings will work fine right out of the box, and you'll be up and running in no time.

In this chapter, we'll take a look at the steps involved in setting up a wireless connection on your home network. We'll focus on some of the main wireless configuration settings you may need to adjust. (Keep in mind that this article will provide you only with general information. You should always refer to the product documentation for specific details on configuring your device.)

Location, location, location

Once you've decided to go wireless, you need to spend some time considering where in your home you'll actually put the Wireless Access Point (WAP). A key element in getting your wireless connection to work effectively is to choose an optimal physical location for your WAP. Of course, the best possible connection between a WAP and a wireless-enabled computer will be within a line of sight. However in a typical home, that's not always possible. Don't worry—a wireless connection can function through walls and floors, allowing you to maintain network access in just about every room in your home. You should even be able to get a wireless connection outside your home within a reasonable distance, allowing you to check e-mail out on the deck or even in that backyard hammock.

To get the best possible coverage inside and outside your home, keep in mind that radio waves emanating from a WAP travel outward in a circular pattern. So, your reception will be better below the WAP's physical location than above it. This means that if you have a multi-story house, you'll get better reception if the WAP is on the same floor or upstairs from you than if it is downstairs from you.

In addition to height, you'll want to choose a centrally located room in your house as your WAP's location. As we mentioned, the higher the WAP the better. Placing the access port on top of a bookcase or other high shelf will provide the best coverage possible.

Also keep in mind that you should avoid placing the WAP near any large metal objects or appliances, such as refrigerators. Brick walls or walls containing a lot of wiring can also cause interference problems.

You may not always be able to immediately identify the best possible location in your home in which to place the WAP. In that case, you may have to do some experimentation. Set up the WAP in what you consider to be the best location and use the wireless network for a while. If you don't get as reliable a connection as you would like, try moving the WAP to another location.

IF YOU CAN'T STAND THE HEAT...

While not a widely known fact, we've discovered that running an average microwave oven while you're using a wireless connection can cause tremendous interference and even completely block a wireless connection. Of course, this problem depends on the location of the microwave oven in relation to the WAP and the wireless-enabled computer. If the microwave oven is between the two devices, you can expect temporary interference while the oven is on. To be on the safe side, make sure that you're not downloading an important file at the same time you decide to make popcorn.

Installing a standalone WAP

If you already have an existing home network and want to add wireless capabilities to it, you'll just need to add a WAP. How you go about doing so will depend on how your current network is laid out.

If you're using a broadband router, which also acts as a hub, your network layout will be similar to the one shown in **Figure A**. If you have a hub connected to a broadband router, your network layout will be similar to the one shown in **Figure B**.

As you can see, regardless of the actual devices involved, you'll connect the WAP directly

to a hub, just like you would if you were connecting a PC to the network. Furthermore, you'll use a standard network patch cable to do so.

Installing a broadband router WAP combination

If you're building a new home network that will include a wireless connection, you'll want to use a broadband router/WAP combination, in which case your network layout will be similar to the one shown in **Figure C**. As you can see, this is a much cleaner setup because you have fewer cables and connections to contend with.

Once you finish setting up the router, you can work on getting the WAP portion of the device up and running.

Installing wireless network card drivers


After you install your WAP on your network, your next step is physically connecting the wireless network card to your computer and installing the necessary drivers. All wireless networking cards come with an installation program that installs the correct set of drivers for your specific operating system and configures the card for use with the WAP. In addition, you'll get a configuration utility that will let you adjust how the wireless network card communicates with the WAP.

If you've purchased a wireless network card and WAP from the same manufacturer, chances are good that the default configuration settings for the wireless network card are in synch with the WAP. As soon as you install the wireless network card and drivers, you should be able to begin using your wireless network immediately.

Wireless network configuration settings


All WAP products will give you a way to configure the WAP. With some products it will be in the form of software that you install on a computer on your network and then use to configure the WAP either across the network or via a USB cable. Still other devices, particularly those that double as broadband routers, can be configured via a Web browser.

Figure A



Add a WAP to an existing home network by connecting it to a broadband router.

Figure B



Add a WAP to an existing home network by connecting it to an external network hub/switch.

DEALING WITH PLACEMENT PROBLEMS

Even though you're building a home network from scratch, you may want to consider getting a separate broadband router and WAP in case you need to position the WAP far away from the broadband router to get the best possible coverage in your home. For example, if your broadband connection enters the home in the basement, that's probably where you'll have the broadband modem and broadband router connected. But it won't be an ideal

place to have the broadband router/WAP combination.

In this case, you could run a regular network cable from your hub in the basement to a more central location upstairs and connect the WAP there. Keep in mind that you can purchase pre-made network cables in 50-foot and 100-foot lengths. If you decide to build your own cable, remember that the maximum distance that you can run a cable is 100 meters, or 328 feet.

As we mentioned earlier, when the wireless network card and WAP are from the same manufacturer, chances are good that you can begin using your wireless network immediately. However, you may want to change some settings, depending on your situation. Also keep in mind that changing some of the settings on your WAP will require similar setting changes to your wireless network card.

With this in mind, let's take a look at some of the most common settings that you may want to alter. Again, remember that we'll be discussing the settings in general and that different products will have different settings. You should always refer to product documentation for specific details on configuring your device.

The Service Set Identifier

The Service Set Identifier (SSID) basically is a name that's assigned to your wireless network. The SSID is much like the workgroup name that you've assigned to your Windows-based network. To communicate, your WAP and all your wireless network cards have to be configured to use the same SSID.

By default the WAP and the wireless network cards will be configured to use a generic SSID. For example, LinkSys wireless devices use *linksys* as the default SSID; Belkin wireless devices default to a SSID of *WLAN*.

As a low-level security measure, you should rename the default SSID to something unique. In many cases, there will be a set limit to the number of characters that you can use for the SSID.


The channel setting

While the 802.11b/WiFi standard supports up to 14 channels, channel numbers are limited by local radio frequency regulations. In the United States, the FCC limits the number of channels to 11. If you wish, you can easily change the channel to any number between 1 and 11. Changing the default channel may help solve interference problems and double as a low-level security measure.

Transfer rate

The maximum transfer rate of an 802.11b/WiFi wireless network is 11Mbps within the specified range, which is typically 100 feet indoors and 500 feet outdoors. When you move beyond that range, the 802.11b/WiFi wireless network standard specifies that the transfer rate will fall back or drop down to

Figure C



Use a broadband router WAP combination.

the transfer rates in **Table A**. Ranges represent typical measurements, and although some manufactures will list varying ranges, the transfer rate numbers will remain the same.

All WAPs and wireless network cards are configured to regularly test distance and automatically select the best possible transfer rate. As you can imagine, running these tests on a regular basis adds overhead to your wireless network and can slow performance. If the computer with the wireless network card is a desktop that never moves or a laptop that you rarely move outside of a specific range, you may want to lock in on a specific transfer rate. Doing so will eliminate the unnecessary testing and associated overhead.

WAP name

On some WAPs, the configuration software will allow you to assign a name to the WAP. This setting is really useful only in a situation in which you have multiple WAPs in the same building and need to keep track of them. In a typical home network setup, naming the WAP is a neat little trick, but not necessary.

Security settings

The batch of settings that may really deserve your attention has to do with security. Why would you want to change the security settings? Well, when you use a wireless network, your data is being transmitted through the air via radio waves. That means that anyone who is within range and using a computer with a wireless network card can potentially hack into

your network and either steal your data or simply piggyback off your Internet connection. Keep in mind that anyone attempting this will have to be highly skilled and extremely determined. Furthermore, they would have to be within the range of your WAP.

Now, in a typical neighborhood this really shouldn't be a problem; however, if you live in an apartment building or condominium or are setting up a wireless network in an office building where others may also be running wireless networks within range of yours, you might want to consider taking some precautionary security measures.

Wireless Equivalent Privacy

The first security measure you might want to consider is called Wireless Equivalent Privacy (WEP), which is an encryption scheme. When you enable WEP on your wireless network, data flowing through the air between your WAP and a computer with a wireless network card via radio waves is encrypted before it leaves the sender and then decrypted by the receiver. Of course, this means that both the WAP and the computer with a wireless network card must be configured to use the same level of encryption as well as the same keys that are used to encrypt and decrypt the data.

There are two levels of WEP encryption that you can use on an 802.11b/WiFi wireless network—64-bit and 128-bit. As you can imagine, 128-bit is more secure than 64-bit.

Keep in mind that using WEP will slow down the rate at which your wireless network can send and receive data. This is due to the fact that the encryption and decryption process takes time. And, of course, using 128-bit encryption will take more time to perform than 64-bit encryption.

WEP can generate the encryption keys either automatically or manually. Doing so manually is a complex procedure, so most often you'll want to use the automatic method. In this procedure you'll simply enter a password called a *passphrase*. When you do so, the software will automatically generate the encryption keys used to scramble and unscramble the data. Of course, this means that you must enter the exact same passphrase

Table A: *Transfer rates*

Location	Range in feet	Transfer rate in Mbps
Indoors	100	11
	165	5.5
	230	2
	300	1
Outdoors	500	11
	885	5.5
	1300	2
	1500	1


on the WAP and all computers with a wireless network card.

MAC Address Filtering

The first security measure you might want to consider is called MAC Address Filtering, which is a more subtle security scheme than WEP. In this scheme, you'll gather and create a list of all the MAC addresses assigned to each wireless network card in your wireless network. You'll then enable the MAC Address Filtering feature on your WAP and enter each of the MAC addresses from your list into the filter. When you do, only those computers whose

wireless network card has a MAC address on the WAP's list can access the wireless network.

Restoring the default settings

If the idea of messing around with settings for your WAP and wireless network card makes you uneasy, don't worry. Almost all configuration software contains a setting for restoring the defaults. So experiment as much as you want, and if you ever feel that you've messed something up, just click the Restore button. 

Diagnosing wireless network performance problems

Aug 5, 2003

By Brien M. Posey, MCSE

Low-budget Wi-Fi networks are extremely popular today, but they are not the only types of wireless networks in existence. There are actually dozens of types of wireless networks ranging in price from under a hundred dollars to millions of dollars. While Wi-Fi problems might not be a big deal to correct, it is a huge problem if you have just spent half a million dollars on a wireless network and the network doesn't perform as expected. In order to get peak performance out of your wireless network, you need to know some common causes of poor performance on both Wi-Fi and non-Wi-Fi wireless networks.

Too many devices

One of the most common problems with wireless networks is having too many wireless devices within close proximity. This problem can be easily avoided by obtaining a professional site survey prior to installing any wireless equipment. Unfortunately, it seems that a

lot of networking people think that if they can install a Wi-Fi network, then the rules are the same for other wireless networks; so, they try to install the devices themselves.

The problem is that certain types of wireless devices are very particular about how many devices can be located within an area. A DS-11, Direct Spread Spectrum network is a good example of this. DS-11 networks have a total of 11 available channels. Because of this, it might stand to reason that you could use 11 different networks within an area without interfering, as long as the networks were on different channels. This isn't the case, though.

The entire concept behind a spread spectrum network is that multiple channels are used in an effort to boost available bandwidth and to increase security. In a DS-11 network, it is only possible to colocate up to three wireless networks before the frequencies start interfering with each other, because each DS-11 device is using multiple channels.

DS-11 isn't the only technology that's subject to this limitation though. You may have heard of a wireless networking technology called FHSS (frequency hopping spread spectrum). FHSS is a type of spread spectrum similar to DS-11 that operates in the 2.4- to 2.483-GHz frequency range. Within this range, there are 79 individual channels and 78 different frequency-hopping sequences that may be used. Even with so many available channels, this type of networking is limited to 15 co-located networks.

The lesson here is that it is extremely important to have a professional site survey conducted prior to installing any wireless hardware. A professional site survey will tell you if other devices already exist within the area that might interfere with the network that you plan to install.

Line-of-sight networks

One commonly used type of wireless networking is line-of-sight. Line-of-sight networks may use either radio signals or lasers to transmit data between two points. As the name implies, line-of-sight technology requires that the transmitter and the receiver have a clear line-of-sight between each other.

There are lots of different problems that can occur with line-of-sight networks. For starters, the antennas used for line-of-sight networks can be notoriously difficult to align. You can easily take care of this problem, though, by having the network professionally installed or by purchasing hardware with a self-aligning mechanism.

I personally like the self-aligning mechanism for more reasons than just that of easy installation. Line-of-sight networks are often used to beam signals between two buildings. The problem is that tall buildings tend to sway a little bit on windy days. Although the swaying may not be more than a few inches in either direction, this is often enough movement to disrupt a wireless network signal. Self-aligning hardware can keep the antennas correctly positioned even if the buildings move.

Keep in mind that although buildings are the most notorious for moving, towers can also move. About nine or ten years ago, I was

experimenting with satellite Internet access. I bought a dish in another state. After making the long drive home, I realized that I didn't have clear line-of-sight with the satellite. To compensate for this, a friend came over and helped me to construct an aluminum tower that was just over 20 feet tall.

For the first couple of days, the signal worked great, but then there was a windy day. Although the naked eye couldn't pick up on much swaying, the signal faded as the tower moved back and forth in the wind.

Maintaining alignment between antennas on a line-of-sight network is important, but it's only half of the battle. It is also important to take the fresnel zone into consideration.

Imagine for a moment that you are standing at one end of a field with a large flashlight trying to illuminate a target at the other end of the field. The beam of light at the far end of the field will be much wider than the flashlight. This illustrates a principle of light called divergence, in which light spreads out as it travels.

While laser light does not have nearly as high of a divergence rate as that from a flashlight, it does exist. Radio signals are also subject to this phenomenon and spread out as they travel.

The problem with line-of-sight networks is that a lot of people don't take divergence into effect. I have seen too many people look out a window, and—if they can see the target—assume that they have a clear line-of-sight to it. However, as the signal spreads out, signal strength is reduced. If you want to receive the signal at full strength, it is important that the receiving antenna have a clear line-of-sight to the entire inbound signal, not just a part of it. The area encompassing the signal is called the fresnel zone.

The fresnel zone identifies the area making up the signal. If there is an object that partially obscures the signal, part of the signal strength would be lost because of the object in the middle. Most line-of-sight networks lack the signal strength to penetrate such obstacles. This shows the importance of a good site survey.

Improper equipment

Another cause of wireless network problems is the use of improper equipment. Earlier, I said that there are dozens of different types of wireless networks. One of the reasons why there are so many different types is because different installations have different requirements. Things like desired bandwidth, climate, distance, and obstacles all play a part in the equipment requirements. If you choose the wrong equipment, your network simply won't perform adequately.

One of the main pieces of equipment that you need to take into consideration is your antenna. Even if you have purchased the correct radios and have done a good job planning your network, poor antenna choices will undermine all of your efforts. While this article isn't intended to be a comprehensive guide to choosing an antenna, I want to take a moment to discuss a few of the more common antenna types, just to give you an idea of why antenna choice is so important.

One of the most common types of antennas is the parabolic dish. This antenna looks like a satellite dish and is commonly used for line-of-sight applications. Most of the time, networks using this type of antenna lack the signal strength to penetrate obstacles but can communicate at great distances. A variation of this type of antenna is the parabolic grid. A parabolic grid works similarly to a parabolic dish but is better suited to windy environments.

Another type of antenna is a panel or sector antenna. This antenna functions like a parabolic dish, but it looks more like a pizza box. These antennas can accept signals varying from 60 to 180 degrees and are suitable for wide area broadcasts.

Still another common type of antenna is the omni. An omni looks like a CB antenna or like a radio antenna that would be used on a boat. An omni has a 360-degree coverage area, but only along a flat horizon. This means that the signal will travel out in all directions, but it won't really travel up or down.

An alternative to an omni is a patch antenna. A patch antenna is a small circular antenna that also has a 360-degree coverage area. Unlike the omni though, a patch antenna

does not have a completely flat horizon. Patch antennas are used primarily for indoor networks.

Poor antenna connection

Yet another common problem is poor antenna connections. A wireless signal is at its strongest when it leaves the receiver. However, there is usually a barrel connector that connects the antenna cable to the receiver, and another barrel connector that connects the antenna cable to the antenna. Barrel connectors diminish the signal strength greatly, as does the length of the antenna cable and even the antenna itself.

Your goal should be to minimize signal loss. To do so, don't use any more barrel connectors than are absolutely necessary, and use the minimum practical cable length. I should also point out that using amplifiers is usually a bad idea. Amplifiers not only amplify the signal, they also amplify noise. More importantly, though, they generally require you to use more cable and a couple more barrel connectors than would be required to connect the radio directly to an antenna, thus diminishing the signal quality. Even if signal distortion were not an issue, an amplified signal often exceeds FCC-mandated signal strengths.

Recently, a friend told me about a network in which the owner was having problems with poor signal strength. The radio was linked to an amplifier. On the other side of the amplifier was a splitter and two antenna cables, which fed two large antennas. Because of the resistance of the three cables, six barrel connectors, two antennas, and the splitter, virtually no signal was being produced. In this situation, my friend simply connected the radio directly to an antenna via a single cable and two barrel connectors and the radio began to perform as it was designed to.

Wi-Fi performance problems

Since Wi-Fi networks are so popular, I wanted to take a moment to discuss some of the problems that are common to Wi-Fi network performance. Although there are some long-distance Wi-Fi implementations in existence, Wi-Fi is designed primarily to be an indoor networking solution. Therefore, this section will address performance problems in an indoor environment.

Just as the most common performance-related problem on big, expensive, outdoor networks is poor signal strength, Wi-Fi networks tend to suffer from poor signal strengths as well. Access point type and position is very important. Some Wi-Fi implementations simply work better than others around obstacles.

For example, several years ago, I bought an 802.11B wireless access point. After installing this access point, I was able to get a wireless signal anywhere in my entire house or yard. Although I liked the convenience, I was always frustrated by the slow speed.

When 802.11A became available, I installed an 802.11A network in my home. This network operates on a frequency of 5.8 GHz as opposed to the 2.4-GHz frequency used by 802.11B. This means that data rates are much higher. However, a 5.8-GHz signal has much more trouble penetrating obstacles than a 2.4-GHz signal does. Consequently, I now have a very fast wireless network, but there are places in my home where I simply can't get a signal.

Another common problem with Wi-Fi networks is that an access point may become oversubscribed. For example, the access point that I'm using in my home supports up to 256 simultaneous Wi-Fi connections. Even the first access point that I ever bought back in 1999 could accept up to 64 connections. The problem is that these high numbers of wireless connections tend to be impractical. It has been my experience that performance starts dropping off once more than about 10 clients are using a single access point.

It may seem that the solution is to add more access points so as to reduce the workload on existing access points. While this is a solution in some situations, having too many access points can cause problems because of interference. If you do have more than one access point, it's a good idea to lock each NIC to a specific access point.

The reason for this is that wireless NICs are designed to roam from one access point to another and to latch onto the access point that has the strongest signal strength. If multiple access points exist in an environment, and two or more have comparable signal strengths, then a NIC may constantly switch back and

forth between access points. This greatly reduces network performance.


Locking a NIC to a specific access point accomplishes three things. First, it eliminates the constant switching between access points. Second, it increases security because no NIC may use an access point unless you have specifically authorized it to do so. Finally, it prevents any of your access points from being oversubscribed.

One last problem that tends to occur with wireless NICs is something called multipath. The best way to describe multipath is by comparing it to a television. You have probably seen a television that used rabbit ears or another type of air antenna. With air antennas, it's common to have some channels in which the main image on the screen is superimposed with a ghosted image. The ghosted image is caused by multipath.

Multipath is caused when a signal bounces off of nearby objects and arrives at the receiver at different times. For example, suppose that you place a wireless NIC into a room that also contained an access point. As the access point transmits a signal, the signal spreads out in all directions and finds its way to the NIC. However, the signal may also bounce off of another object in the room and find its way to the wireless NIC by this diverted path. This means that the wireless NIC is actually receiving the signal twice.

In the real world, it's almost impossible to get rid of multipath signals. The most that you can do is to try to reduce it by not having any large metal objects in the proximity of the access point or the wireless NIC.

No wires doesn't mean no problems

As you can see, there are many different types of wireless networking problems. Over the past year, I have seen a lot of networking professionals make a trip to the computer store and spend a couple hundred dollars on a wireless access point and a few NICs. Often, when this equipment gets installed the results can be disappointing. Once you know what can cause problems with wireless networks, you can figure out how to solve those problems and get the most out of your investment. 

Fix hardware and configuration issues common to wireless LANs

Aug 20, 2002

By Brien M. Posey, MCSE

With decreasing prices of wireless hardware, wireless networks are fast becoming more popular in small office networks. Both the cost savings and the ease of using wireless LANs are beneficial to the small office—until something goes wrong. Then it becomes all too apparent that, while wireless networks are growing, troubleshooting resources for wireless LANs are not.

When a wireless network fails, there are a few key areas to look to first. In this article, I'll discuss some of the more common hardware problems that can cause a wireless network to fail. As well, I'll cover the configuration issues that can also plague a wireless LAN. With this information, you can troubleshoot your wireless network with confidence. (This article assumes that you're troubleshooting an infrastructure network and not an ad hoc network.)

Hardware troubleshooting

When you have only one access point and only one wireless client with connection issues, then you've already determined the scope of the problem. It's your one client that is having trouble attaching to the network. However, if you've got a larger network, then the process of determining the scope of the problem becomes a little more involved.

If lots of users are having trouble connecting, but there are still some users who are able to work, the problem is most likely that your network has multiple access points and that one of the access points is malfunctioning. Often, you can take an educated guess as to which access point is malfunctioning by looking at the physical locations of the users who are having the problem and then figuring out which access point serves that portion of the building.

If no one can connect to the wireless network, then there are several things that could be going on. If your network uses a single

wireless access point, one possibility is that the access point is malfunctioning or contains a configuration error. The problem could also be related to radio interference or to a break in the physical link between the wireless access point and the wired network.

Check connectivity to the access point

First, you should perform a communications test to see if the access point is responding. To do so, open a Command Prompt window on a PC on your wired network and ping your wireless access point's IP address. The wireless access point should respond to the ping. If it doesn't, there's either a break in the communications link or the access point is completely malfunctioning.

To figure out which is the case, try pinging the access point's IP address from a wireless client. If the wireless client is able to ping the access point successfully, then the problem is almost certainly a broken communications link, such as a damaged cable.

If the wireless client is unable to ping the access point, then the access point could be malfunctioning. Try unplugging the access point to reset it and then plug it in again. Wait for about five minutes and then try pinging the access point from both the wireless and the wired clients again.

If both pings still fail, then it is likely that the access point is damaged or has an invalid configuration. At this point, I recommend focusing your initial efforts on getting the access point to communicate with the wired network. Plug the access point into a known-good network jack using a known-working patch cable. You should also verify the access point's TCP/IP configuration. After doing so, try pinging the device from a wired client again. If the ping still fails, then the unit has probably been damaged and should be replaced.

Configuration issues

I've found that wireless networking equipment is fairly reliable, and the vast majority of

problems are related to the network's configuration rather than to a hardware malfunction. With this in mind, I'll discuss several common hardware configuration problems that lead to a disruption of wireless services.

Test the signal strength

If you can ping the wireless access point from a wired client but not from a wireless client, then the access point is probably just experiencing a temporary problem. If the access point continues to have problems, I recommend checking the signal strength. Unfortunately, there's no standard method for doing this. Most wireless NIC manufacturers, however, include some mechanism with the NIC for measuring signal strength.

Try changing channels

If you determine that you're getting a weak signal but nothing has physically changed in your office, then I recommend attempting to change channels on the access point and on one wireless client to see if a different channel improves the signal strength. I run a wireless network in my home office, and I've found that one of my cordless phones interferes with my wireless network when the phone is in use. 802.11b wireless networks function on the 2.4-GHz frequency, just like many higher-end cordless phones. Changing channels on all of your wireless clients can be a big undertaking. Therefore, I recommend testing the new channel with one client first. Remember that your problem could go away as soon as someone hangs up a phone or turns off a microwave oven.

Verify the SSID

A while back, I took my laptop to a friend's house to work. Because my friend had a wireless network in place, I decided to connect to his network for the duration of my visit. Upon returning home, I didn't use my laptop for a couple of weeks. The next time that I went to use my laptop, it wouldn't connect to my network. The problem was that I had forgotten to reset the SSID (Service Set Identifier) back to my own network identifier. Remember, if the SSID doesn't specify the correct network, then you won't be able to ping the access point.

Instead, your laptop will ignore the access point's existence and search for an access point with the specified SSID.

Verify the WEP key

I recommend checking out the wired equivalent privacy (WEP) encryption configuration next. If WEP is configured incorrectly, you will not be able to ping the access point from a wireless client. Different brands of NICs and access points require you to specify the WEP encryption key differently. For example, one brand requires you to enter the encryption key in hex format, while another brand requires the key to be entered in decimal format. Likewise, some brands support 40-bit and 64-bit encryption, while other brands support only 128-bit encryption.

In order for WEP to function, all settings must match exactly between the client and the access point. I have run into several situations in which clients that seemed to be configured perfectly simply could not communicate with an access point that was using WEP. During these situations, I usually had to reset the access point to the factory defaults and reenter the WEP configuration information. Only then did WEP begin to function.

Tricky WEP configuration issues

By far the most common configuration-related problems involve the use of the WEP protocol, so WEP deserves some more discussion. Troubleshooting a WEP problem can be especially tricky, because a WEP mismatch has symptoms that are similar to a more serious failure. For example, if WEP is configured incorrectly, a wireless client won't be able to get an IP address from a DHCP server (even if the access point has a built-in DHCP server). If the wireless client is configured to use static IP addresses, the wireless client won't even be able to ping the access point's IP address, thus giving the illusion that no connection exists.

The trick to figuring out whether a problem is related to a WEP configuration error rather than a hardware malfunction is to be aware of the diagnostic capabilities built in to the NIC driver and the operating system. For example, one of my laptops is running Windows XP and has a Linksys wireless NIC. Notice in

Figure A that if I move my mouse pointer over the top of the wireless icon in the taskbar, I see a summary of my connection information. In this case, the connection strength is Excellent. As long as the channel and SSID are configured correctly, you can connect to the access point, even with a WEP configuration error. Had there been a physical connection problem, the connection strength would be None, not Excellent. Linksys cards will show you the connection strength whether WEP is configured correctly or not. Therefore, you can validate that a connection exists, even if you can't ping the access point.

If you right-click on the wireless networking icon in the taskbar and select the View Available Wireless Networks command from the resulting menu, you'll see the connect to Wireless Network dialog box. This dialog box displays the SSID of any wireless network on your present channel to which you are not currently connected. If the name of your wireless network shows up on this list but you can't seem to connect, then you can rest assured that your connection is good and that you've got a configuration problem.

NOTE

An interesting side note is that the Connect To Wireless Network dialog box also includes a field where you can enter a WEP key when you try to connect to a wireless network. There have been times when I absolutely could not connect to a particular wireless network unless I went through this dialog box and manually entered the WEP key. After doing so, the network became available to me.

DHCP configuration issues

Another tricky problem that can prevent you from successfully interacting with a wireless network is a DHCP configuration error. The DHCP server that you connect to can play a major role in whether you are able to use a wireless network.

Many of the newer access points have an integrated DHCP server. Typically, these access points assign the 192.168.0.x address

Figure A



The signal strength is a big clue as to the nature of your problem.

range to clients. Often, DHCP access points will not accept connections from clients to which they have not issued an IP address. This means that clients with static IP addresses or clients that might have somehow acquired an IP address from another DHCP server could be unable to connect to the access point.

The first time that I installed an integrated DHCP server access point onto my network, I decided to allow the access point to assign IP addresses to my wireless clients. However, my network uses the 147.100.x.y address range. This meant that although wireless clients were able to communicate with the access point and were able to acquire an IP address, they were unable to interact with the rest of my network because of the IP address range mismatch.

There are two solutions to this problem:

- ▶ Disable the access point's DHCP services and allow the wireless client to lease an IP address from a normal DHCP server.
- ▶ Override the IP address range by configuring the DHCP address scope with your own block of IP addresses.

Either solution will work, but you'll have to work within the limitations imposed by your access point's firmware. Many access points

will allow you to use only one solution or the other, but not both.

Problems with multiple access points


Suppose for a moment that two access points are in use, both with the default settings. If this is the case, then both access points are assigning clients IP addresses in the 192.168.0.x address range. The problem is that the two access points are completely unaware of which IP addresses the other access point has leased. Therefore, it's only a matter of time before there are duplicate addresses on your network.

The solution to this problem is to define a unique scope of addresses for each access point. By doing so, you'll prevent IP address overlaps.

Watch out for client lists

Some access points contain an allowed client list, which can be the root of wireless configuration problems. The allowed client list is a list of MAC addresses of permitted wireless

clients. This is a security feature that's designed to prevent unauthorized users from connecting to your network. Normally, the allowed address feature is disabled by default. However, if a user has accidentally clicked the Enable button, then the allowed address list will be enabled but won't contain any MAC addresses. This means that no wireless clients will be able to connect to the access point, regardless of any other configuration settings.

I've also seen the allowed address list become a problem when multiple access points are in use. Many administrators incorrectly assume that just because they enter the allowed addresses into the list, the addresses are then globally permitted to access the network. However, in most cases, this simply grants the users permission to access the network through the designated access point. If you want users to be able to go through other access points, you'll usually have to configure those access points separately. 

Troubleshooting the wireless woes

May 7, 2002

By Luke Mason

As if the *usual* things that can and do go wrong in IT aren't enough to drive us crazy, we IT managers have to deal with the occasional anomalies, those annoying little problems that at first appear to have no cause and, therefore, no solution. These are the times when you have to step up and become IT's version of Sherlock Holmes.

I recently encountered an interesting problem on the job. When Mark, an employee, phoned me and said, "Can you come and have a look at my computer? It's gone all funny," I was pretty certain that I was going to have to free up at least half an hour of my afternoon.

Mark is one of those users who are a blessing to the lazy among us and a nightmare to the diligent. He will ignore *any* error message that comes his way, no matter how serious he thinks it sounds. He realized that restarting his PC was a good way of sorting out some errors. When NT presents him with a dialog stating that "a domain controller for this domain could not be contacted," he clicks OK and tries to continue. When Outlook starts whining about address books, he again homes in on the OK button and starts to worry only when he meets some nonsense about POP3 servers, "whatever they are."

The background

Mark's PC is running on our semiexperimental wireless network along with three other people's on the same floor. Pinging anything other than *localhost* proved that there was no connection to the network, and both servers were humming and clicking merrily away with no sign of any problems. Everyone else on the wireless net was still connected, so it had to be Mark's computer. He still had a green signal light on the network card, so *something* was getting through, but I didn't know how much.

The configuration utility that comes with the Netgear MA301 cards that we use helpfully installs itself in the *startup* folder of all users of the machine, and even comes complete with two handy little meters showing signal quality and link strength. Rather unhelpfully, the utility refuses to run unless

you have administrator rights, so it has to be removed from the startup profile again shortly after installation. Logging on as an administrator, I could see that the signal

strength was at 6 percent—not very good when the access point is only a few feet away and 802.11b is supposed to provide full 11 Mbps at a distance of 30.5 metres (100 feet)!

What went wrong

Wireless cards work, in the same way as mobile phones, on a line-of-sight principle. Anything solid that gets in their way absorbs some or all of the signal. That's why mobile phones often work better if you stand next to a window; glass doesn't absorb as much of the radiation as a concrete or brick wall. A desk had recently been moved into the line between the access point in my office and Mark's PC. I hadn't expected a few sheets of pine to interfere with the signal, but this was the only possible cause of the dropped connection. This problem only manifested itself a few days afterward, and why the signal didn't drop as soon as the desk was moved, I couldn't fathom. I didn't have a spare card around to swap into Mark's PC, so the only thing I could try was to move the access point.

The solution

With the wireless network's newfound sensitivity to carelessly placed inanimate objects, I had to be careful not to block someone else out of the signal. In true "hit it and hope" fashion, I wiggled the antennae on the access point and moved it as far as its patch lead would allow: the 6 percent meter didn't change. The only solution I could think of was to run a longer CAT-5 cable to the other side of the room and place the access point on top of my cabinet. With this done, Mark's signal strength shot up to 97 percent, and his PC could once again join in the fun on our domain.

Why did this happen? I'm not sure. I've walked a laptop with a wireless card down two flights of stairs to the second floor, and I still had a signal higher than 6 percent. Two users

sit on the other side of a wall to the access point and receive signals in the high 90s. I can't really believe that a wooden desk could block that strong a signal; after all, that would make wireless net-

Wireless cards work on a line-of-sight principle. Anything solid that gets in their way absorbs some or all of the signal.

works in general a little ineffective, wouldn't it? One guess is that the wireless card has a fault on it. Possibly it affected signal strength, and moving the access point fixed it, or maybe it was simply an intermittent problem that corrected itself. Maybe the desk or something in it had strange properties in blocking microwave radiation. Perhaps Mark's mobile phone interfered with the signal, equivalent to your car stereo's announcing an incoming call by barking and groaning at you through its speakers.


The point is that when you're dealing with wireless networks, you can't ever really be sure. At least with CAT-5 you can safely assume that the cabling within your building is sound. Swapping a patch cable is simple and, unlike wireless cards, they don't cost \$100, so you can afford to keep spares. Once the integrity of the physical link is established, you can start the process of examining the software configuration of the PC. But with a wireless network you're feeling your way in the dark. I don't mind keeping track of a

few cards, but I wouldn't like to go anywhere near an office with more than 10.

An IT head-slapper

On the subject of network devices and wireless networks, how safe are your power sockets? A strange question, I know, but bear with me. We recently hired a new cleaning firm at my company, and one of the new recruits apparently took silent and stealthy pleasure in unplugging our Cisco router so that they could plug in a vacuum cleaner—leaving our mail server, the online backup, and anyone trying to connect wondering what had happened. Never mind that the router was mounted on the wall, below what I thought was a pretty terrifying patch panel stuffed with blinking lights and sprouting yellow and blue patch cables. Never mind that the voicemail system was purring

and ticking away next to it in what even I consider to be an unnecessarily intimidating fashion. The fact that the plug was in the wall in the normal way was permission enough for the cleaner to disconnect it. And it just served to show me that there are a lot of little hazard areas that I'm unlikely to think about until danger strikes.

There's really no foolproof solution for this problem. I doubt many people can afford to UPS their routers and switches, and if your patch panel doesn't have power built in, you're left with little alternative than to use the power sockets on the wall. The best bet? Place a clear sign on the wall. Now go and write "Do Not Unplug" on all of your network equipment, before it happens to you. 

Troubleshoot wireless networking antennas

Nov 7, 2002

By Scott Lowe, MCSE

When troubleshooting a wireless network problem, it's important to have a thorough understanding of the technology and your options for fixing the problem. In this article, I'll show you how to troubleshoot signal problems on wireless networks, focusing on potential problems with antennas.

A FEW ASSUMPTIONS

As Sherlock Holmes said: "When you have eliminated the impossible, whatever remains, however improbable, must be the truth." This holds true when troubleshooting network problems. For the purposes of this article, I'm assuming that you've taken troubleshooting steps at the access point and physical network and that the problem has been narrowed down to the antenna or connecting hardware.

Get the pieces right to start with

Understanding the components that make up a wireless network is critical to network support and troubleshooting. One of the primary components of a wireless network is the antenna. The type of antenna you choose directly affects its performance as part of the network, as well as the application for which it's suitable.

Wireless networking commonly uses two types of applications. The first is a site-to-site application in which two physically separate sites or buildings are connected to each other using a wireless bridge. The second type is client-based. A wireless access point is deployed to directly support laptops or other wireless clients for network connectivity. Each of these two types of applications has antennas that are better suited to it than the other.

Four basic types of antennas are commonly used in 802.11 wireless networking applications: parabolic grid, yagi, dipole, and vertical.

Parabolic grid

Perhaps the most powerful antenna for site-to-site applications is the parabolic grid antenna. A parabolic grid antenna can take many forms, ranging from something that looks like a satellite TV dish to one that has the same shape but is made of a wire grid instead of having a solid central core. This type of antenna is a unidirectional antenna, meaning that it transmits in one specific direction: the direction at which the antenna is pointed. **Figure A** depicts a parabolic grid antenna.

Yagi

A yagi antenna is slightly less powerful than a parabolic grid, and it's suitable for site-to-site applications at lesser distances than a parabolic grid. Like the parabolic, the yagi is also a unidirectional unit. A yagi antenna consists of a series of metal spokes radiating from a central core. The whole thing is covered by a tubular plastic housing called a radome, so you seldom see the actual antenna elements. **Figure B** depicts a yagi antenna with a cutout showing what the internal elements look like.


Dipole

A dipole is a bidirectional antenna, and its radiation pattern extends in two directions outward, as shown in **Figure C**. It generally consists of a base with two antenna spokes going in opposite directions. You'd generally use a dipole antenna to support client connections rather than site-to-site applications.

Vertical


A vertical antenna is exactly what it sounds like: an antenna that sticks in the air. A vertical antenna's radiation pattern extends in all directions from the unit, losing power as the distance increases, as shown in **Figure D**. Like the dipole, you'd primarily use a vertical antenna for client support. Most wireless base stations come with a small vertical antenna. A vertical antenna is omnidirectional, meaning that the signal radiates in all directions.

Figure A




A parabolic grid or dish antenna and its radiation pattern

Figure B



A yagi antenna and its radiation pattern

Figure C



A dipole antenna and its radiation pattern

Antenna specifications

Understanding the different antenna types is only the beginning. Each antenna type has a number of specifications that directly affect how well it works. These specifications are antenna gain, beam width, loss, and radiation pattern.

Antenna gain

This is a measurement of how well the antenna focuses a signal. This is typically measured in dBi (decibels relative to isotropic radiator—a theoretically “perfect” antenna) and is based on decibels, which is a logarithmic measure of relative power. The dBi is computed by comparing the output of the antenna to a theoretical isotropic radiator (antenna) with a dBi of 0: the higher the dBi measurement, the higher the power level of the antenna.

Beam width

The beam width is the area radiating outward from the antenna where the signal within a specific angular distance is above the “half power” of the peak intensity of the antenna. The beam width is also loosely used to determine the antenna type. A parabolic grid antenna is a unidirectional antenna with a very low beam width, which means that it needs to be very carefully aimed at its partner in order to be effective. A vertical, omnidirectional antenna has a very high horizontal beam

width, which is why it’s suitable for roaming client connections; however, its vertical beam width will be lower. In general, there’s an inverse correlation between beam width and antenna gain, which means that the required accuracy for aligning antenna goes up as the gain increases because the beam width decreases.

Loss

Loss is an important factor when deploying a wireless network, especially at higher power levels. Loss occurs as a result of the signal traveling between the wireless base unit and the antenna. Since these units are always connected by a cable, there will always be loss. You can minimize loss by using the appropriate type of cable in the minimum length required to make the connection.

Radiation pattern

In Figures A through D, I showed you a sample radiation pattern for each type of antenna. Every antenna has a unique radiation pattern determined by its construction. This radiation pattern is a three-dimensional radiation field of the antenna’s output. Some antenna manufacturers supply sample radiation pattern specifications for their equipment.

You can use these specifications to determine how far the signal from a particular antenna can travel before becoming unusable. As a rule of thumb, a directional antenna has a conical pattern of coverage that radiates in the direction that the antenna is pointed, while an omnidirectional antenna’s area of coverage is shaped like a doughnut.


Troubleshooting some common problems

A good understanding of wireless networking antennas makes troubleshooting much easier. Exactly how you solve the problem depends on the type of connection you’re trying to make—site-to-site or local wireless connections.

Troubleshooting site-to-site connections

It’s late October in your beautiful upstate New York town, and you’ve just finished putting up your last wireless antenna. Now you’re going

Figure D



A vertical antenna and its radiation pattern

from site to site, aiming them in the right direction. You spend a few days getting everything just perfect and even get close to 100 percent efficiency on your link! Life couldn't be better. Then, in April, your links start to degrade slowly, and by May, some are almost unusable. You check your antenna connecting hardware and everything looks good. What could have happened? Unfortunately, trees grow leaves, and leaves are a Wi-Fi killer because they contain water. You'd be better off with a brick wall in the way!

The best way to address this problem is to install your hardware at the time of year when conditions will be the most difficult, such as in the late spring when all of the leaves are out and the summer drought hasn't yet begun. If this isn't feasible, I recommend installing your antennas on towers on top of the buildings you're connecting (keep in mind that trees also have a tendency to grow, so plan accordingly). If this isn't possible either, you might consider a higher-power antenna, but keep in mind that the FCC limits the total transmission power to 1 watt, or about 36 dBm.

When problems arise suddenly with a site-to-site connection, it's probable that something has happened to one of the antennas. You may need to physically check your antennas to make sure that none have been damaged, have fallen off the building, or have been bumped around and are now not aimed correctly. If you're using a solid parabolic grid antenna and live in an area where high winds are common, you may want to consider replacing it with a mesh dish in order to prevent potential wind damage.

Other problems with site-to-site connections involve interference from various sources, including other 2.4-GHz installations. Most of today's wireless networking systems use the 2.4-GHz spectrum. If your company and the company next door both decide to implement a wireless network between your various sites, you may notice degradation in performance because of interference caused by your neighbor. You may need to use different channels for your installation.

If that's not possible and you need a performance boost, you may have to migrate to

802.11a technology, which would require you to replace all of your equipment. Equipment using the 802.11a standard operates in the 5- to 6-GHz range, but it's much more expensive than 802.11b equipment, and it's not backward compatible with 802.11b. However, 802.11a gear can operate at speeds of up to 54 Mbps, or close to five times the 11 Mbps limit of 802.11b. If at all possible, I don't recommend ripping out your 802.11b equipment and replacing it with 802.11a. A new standard, 802.11g, has recently been approved. It will allow transmission speeds of up to 54 Mbps in the 2.4-GHz range, and it's compatible with 802.11b.

Troubleshooting local wireless network problems

Because of the "walking around" aspect to client wireless networking and because a wireless network is generally deployed inside buildings, you're more likely to experience unusual problems with these types of connections than in your fixed point-to-point connections.

Some common wireless connectivity problems have to do with your distance from the antenna, as you'd expect. However, there are also certain places that can create problems with wireless connectivity. One such place is directly under a vertical antenna that is pointing upward. As mentioned earlier, an omnidirectional antenna has a doughnut-shaped area of coverage, which means that there's a hole right in the middle. If you're working in the area covered by the hole and you aren't able to connect to the network, try moving your wireless device, moving your base antenna, or mounting the antenna upside-down on the ceiling instead. If you're not using an omnidirectional antenna for your indoor client application, you should replace the directional antenna that you are using. If you're using a directional antenna because of its increased range, just add a second access point to cover the same distance. In the long run, you'll have a more efficient system in place, as well as better throughput from more areas.


Other problems have to do with the way that the wireless adapter fits in the PC. To achieve the most desirable coverage area, the

antenna for a wireless adapter should be pointed up. Unfortunately, most wireless adapter antennas point horizontally, which greatly limits their range. The best way to correct for this is to attempt to point the side of your laptop with the adapter antenna toward the wireless access point antenna. This may solve your connection problem and potentially boost the strength of your signal at the same time.

My last point has to do with antenna placement. If you've positioned your base antenna or your client card antenna near a device generating a 2.4-GHz field, such as a cordless phone, you may experience major interference. Try moving the antenna or the source of the interference. Likewise, try not to place objects such as fish tanks or water coolers in the line

of sight between antennas, since water will refract the signal. Finally, avoid placing antennas near large metal objects, microwaves, and other sources of electromagnetic (EM) interference.

That's all there is to it

Troubleshooting antenna problems and boosting network performance are sciences in and of themselves. If you want to truly understand how the signal is generated and how it travels, there are numerous resources available on the Web. Knowing some potential trouble scenarios and how to deal with them can help you solve problems and increase performance at the same time. 

Notes

Notes

File and Share Permissions

- File-sharing permissions in Windows 2000.....73
- NTFS permissions in Windows 200077
- Combining sharing and NTFS permissions in Windows 200080
- Establish the correct file-sharing permissions in Windows XP83
- Effectively set and troubleshoot NTFS permissions in Windows XP86
- Combining sharing and NTFS permissions in Windows XP90

File-sharing permissions in Windows 2000

Jan 25, 2001

By Faithe Wempen, A+, MOUS 2000 Master

One of the reasons people have preferred Windows NT—and now Windows 2000—to the Windows 9x platform has been the ability to set and manage file permissions more precisely and more conveniently. If you use the NT file system (NTFS), you can set file permissions at the local PC level in addition to the file-sharing permissions of the network environment.

But along with all this additional functionality comes complexity and the potential for all kinds of headaches for the network administrator. One harried manager wants to know why he can't access the data on a colleague's PC that he needs to assemble an important presentation; another can't figure out why the intern from the mailroom was able to browse the files he thought he had secured. More options mean more chances for confusion and user error, and if you don't have a thorough understanding of the various permissions and their relationships, it can be nearly impossible to sort out a permission problem and find a solution.

In this article, I'll review the file and folder permissions in Windows 2000. My next article will cover NTFS permissions in Windows 2000. Once you understand Windows 2000 permissions and how they interact, you should be able to troubleshoot permission issues more quickly as they occur on your network.

Overview

In any Windows network environment (peer-to-peer or server-based), you can set sharing permissions for drives and folders. By default, when you set up a PC on a network, no drives or folders on that PC are shared. The local user of that PC can choose to share entire drives or individual folders on a drive. This type of security is not really that secure, however, because it affects only network access. Local access (that is, someone sitting down at the PC and logging on) is wide open.

For drives formatted with NTFS, you can set NTFS permissions. These can affect drives and folders and individual files, too. NTFS permissions affect local users as well as network users and are based on the permission granted to individual user logons, regardless of from where they are connecting. You also have a much wider variety of permissions to choose from with NTFS permissions, so you can more precisely control the rights being granted.

When sharing permissions and NTFS permissions conflict, the most restrictive of the two wins. For example, if someone has full access to a certain file from NTFS permissions but has no sharing permissions to the folder in which it resides, he or she cannot access the file from the network. The user can, however, physically sit down at the local PC containing the file, log in, and access it, because sharing permissions do not affect local access.

Working with shared folders

Shared folders provide remote access to the files on a PC. Folder sharing is available on drives using all types of partitions: FAT, FAT32, or NTFS. It is also available not only in Windows 2000 but also in Windows NT and Windows 95/98/Me and even the old Windows 3.11 for Workgroups (although in a more rudimentary way in that OS).

To share any folders (or any printers, for that matter) on a Windows 2000 PC, File And Printer Sharing For Microsoft Networks must be installed as a networking component. To check for it, right-click My Network Places and choose Properties. Then right-click Local Area Connection and choose Properties. If File And Printer Sharing For Microsoft Networks does not appear on the list shown in **Figure A**, add it by clicking Install and choosing it from the Services category.

After File And Printer Sharing For Microsoft Networks is in place, you can share individual drives and folders. Do so by right-clicking a drive or folder and choosing Sharing.

The Sharing tab of the Properties dialog box opens.

Sharing is slightly different for drives versus files. With a drive, you might see a default share already set up. These have a \$ following the share name, as in **Figure B**. Such shares are for administrative use only; ordinary users will not be able to see or browse a drive shared in this way on the network. Consequently, if you want to share an entire drive like this on your network, you must create an additional share for it.

To create a new share for a drive, click the New Share button and then fill in the share name, any comment you want to make, and a user limit for concurrent usage (if desired). While you are in the New Share dialog box (see **Figure C**), you can click the Permissions button to specify who will have access to the shared drive, or you can save that for later.

For a folder, the process is more straightforward because there are no default administrative shares. By default, a folder is set to Do

Not Share This Folder. To share it, choose the Share This Folder button and then enter a share name, comment, and user limit.

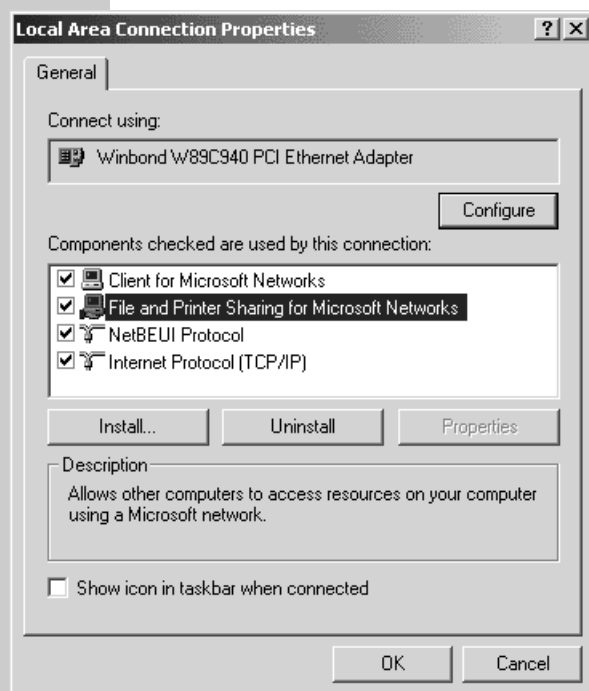
Regardless of whether you are sharing a folder or a drive, you can configure permissions the same way: Display the Sharing tab and click the Permissions button. A Permissions dialog box appears, as in **Figure D**. By default, all permissions are granted to everyone.

If you plan to use NTFS permissions in conjunction with sharing permissions, you might want to leave the sharing permissions set at the default “free-for-all” settings and rely on the NTFS permissions to lock down certain sensitive items. However, if you aren’t going to use NTFS permissions, or if you can’t because the drive is FAT or FAT32, you might want to restrict access at the sharing level.

Note in Figure D that there are three types of sharing permissions:

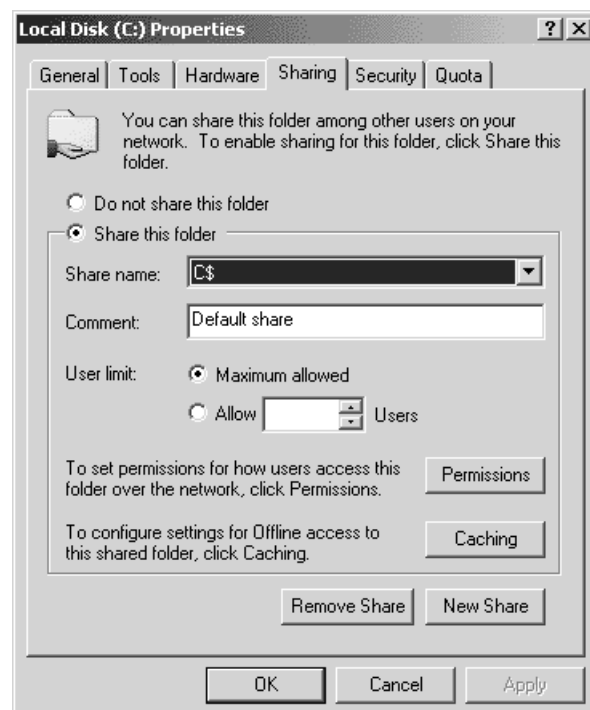
► **Read:** Users can display the contents of the folder, open files, display attributes, and run programs.

Figure A



File And Printer Sharing For Microsoft Networks must be installed in order to share folders over a network.

Figure B



C\$ is the default administrative share for this drive; it does not count as a user-to-user share.

- **Change:** Users have all the rights of Read plus the ability to create new folders and files within the shared folder or drive, open and change files, change file attributes, and delete folders and files.
- **Full Control:** Users have all of the rights of Change plus the ability to take ownership of files and change file permissions.

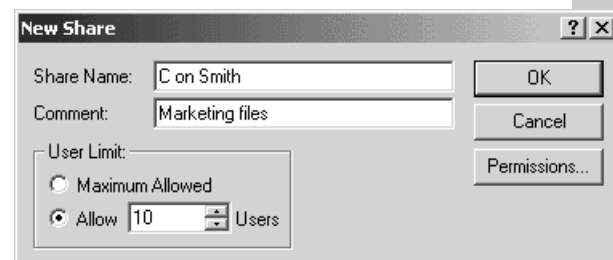
Everything within a shared drive or folder inherits its sharing permissions. For example, if a shared drive has 10 folders, all of those folders have the same sharing permissions as the drive unless they are set otherwise. Permissions are cumulative, which means in the event of a conflict between a specific folder's permissions and those it has inherited from the drive (or parent folder), the most lenient wins. For example, if you allow Read access on a folder but do not allow Change or Full Control on that folder but the drive itself allows Full Control, that folder will also have Full Control access permitted.

For each setting (Read, Change, and Full Control), you can choose the option to Allow or Deny. The default is set to Allow. If you don't want to allow a particular permission,

you simply deselect the Allow checkbox. "Dis-allowing" something (that is, turning off Allow permissions for it) takes away that right but enables the folder to inherit permissions from the parent folder or drive.

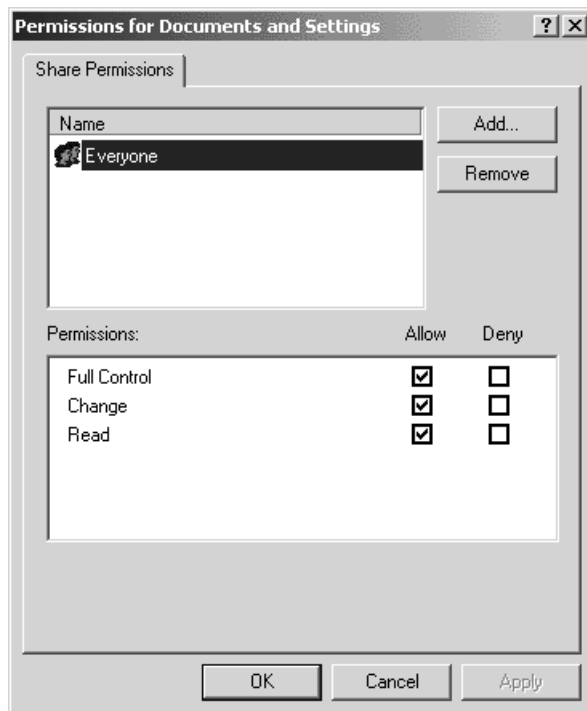
When you share a folder or drive, there is only one group with permissions assigned by default: the Everyone group. That means all users will have the same permission rights to the object, regardless of any group affiliation. You can delete the Everyone group from the list and/or add other groups or individuals to its permissions list. You might, for example, delete the Everyone group from the list entirely or leave it there and set it to allow

Figure C



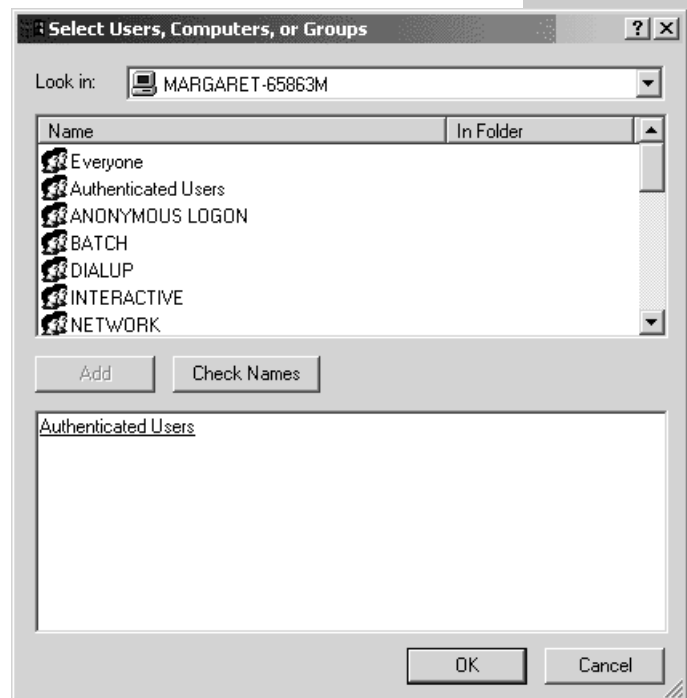
Create a new share to allow other users to access the drive.

Figure D



Limit permission to the folder or drive if desired.

Figure E



Specify other users or groups besides Everyone to receive permissions.

HERE ARE SOME TIPS FOR USING SHARING PERMISSIONS EFFECTIVELY

- ▶ Grant only the permissions that a group or user needs; disallow all others. In most cases, Change permission is all a user needs for a drive or folder. Change enables users to run programs, edit files, and so on.
- ▶ Do not allow Full Control for a drive to the Everyone group. If certain users must have complete control of a drive, assign Full Control to a particular group or create a group for that purpose.
- ▶ Do not use the Deny option unless you have a specific reason to do so. It's easy to forget that you've used the Deny option and spend fruitless hours troubleshooting a file access issue because of it.
- ▶ Assign sharing permissions to groups, rather than individuals, to minimize administrative work.
- ▶ Use descriptive share names to help users locate the shared drives or folders they want.
- ▶ Group folders that need to have the same sharing permissions assigned in a single folder together and then assign the permissions to the parent folder.

TIP: DON'T DENY


The Deny option should be used sparingly, because it overrides any more lenient permissions. For example, if you set Read access for a folder to Deny and the drive on which the folder resides allows Full Control, everything on that drive will have Full Control access except for that folder, which will have no access at all.

Read permission only and then add the Administrators group to the list and grant that group Full Control.

To add a group or user to the permissions list for an object, start from the Permissions

dialog box (Figure D); click the Add button; choose the user or group you want in the Select Users, Computers, Or Groups box (Figure E); and click the Add button. When you're finished, click OK to return to the Permissions dialog box. The users and groups you chose will appear on the Permissions list, ready to have their permission levels set.

Conclusion

In this article, you learned to configure file-sharing permissions for groups and individuals. You learned how permissions are inherited and what happens when file and folder permissions conflict. 

NTFS permissions in Windows 2000

Jan 29, 2001

By Faithe Wempen, A+, MOUS 2000 Master

Setting folder and file permissions gives you some network security, but it doesn't secure your PC desktop. When you use the NT file system (NTFS) in Windows 2000, however, you can set file permissions at the local PC level in addition to the file-sharing permissions of the network environment. In this article, I will cover NTFS permissions in Windows 2000.

NTFS permissions overview

NTFS permissions can be set only on drives partitioned with NTFS. NTFS permissions, like sharing permissions, specify who can access a particular resource, but they work at the local level. That means a user sitting down at a PC is bound by NTFS permissions too, not just a user accessing the resource across a network.

NTFS permissions can be assigned to drives and folders, just like sharing permissions, but they also can be assigned to individual files. Unlike sharing permissions, in which the default setting for a resource is Not Shared, NTFS permissions are set to allow access by default.

Folder and drive permissions

NTFS offers many more types of permission than the simple Read, Change, and Full Control of sharing permissions. For folders and drives, you can assign these permissions:

- ▶ **List Folder Contents:** View folder contents.
- ▶ **Read:** View folder contents, open files, and view file and folder attributes.
- ▶ **Read & Execute:** Same as Read, plus the ability to move through folders to reach other files and folders, even if no permission is granted for those folders.
- ▶ **Write:** Same as Read, plus the ability to create and edit files and subfolders and change attributes.
- ▶ **Modify:** Combination of Read & Execute and Write, plus the ability to delete the folder.

- ▶ **Full Control:** Same as Modify, plus the ability to change permissions, take ownership, and delete subfolders and files.

File level permissions

The permissions for individual files are the same types, except there is no List Folder Contents permission. For files, you can assign these permissions:

- ▶ **Read:** Open the file and view its attributes, ownership, and permissions.
- ▶ **Read & Execute:** Same as Read, plus the ability to run applications.
- ▶ **Write:** Same as Read, plus the ability to change file content and attributes.
- ▶ **Modify:** Same as Write and Read & Execute combined, plus the ability to delete the file.
- ▶ **Full Control:** Same as Modify, plus the ability to change permissions and take ownership.

Just like sharing permissions, NTFS permissions can be set to allow or not, depending on whether the Allow check box is marked. Permissions are cumulative and can be inherited from parent folders or drives. NTFS permissions can also be set to Deny, but use Deny very sparingly.

To set NTFS permissions, you use the Security tab on the Data Properties page for a drive, folder, or file. The controls will seem familiar, as they're almost the same as the ones for setting sharing permissions (**Figure A**).

Inheriting permissions

Notice the check box at the bottom of **Figure A**. When it is turned on, the folder or file will inherit the permissions of the parent object (that is, the drive or folder in which it resides). The gray check boxes in **Figure A** indicate that those permissions are inherited rather than specific to this folder.

If you deselect the Allow Inheritable Permissions From Parent To Propagate To This Object check box, a dialog box appears asking what you want to do about those inherited settings. (You

won't see this on drives, because they have nothing to inherit from, being at the top level already.) You can choose to copy them or to remove them. If you remove them, all permissions and all users that were inherited are stripped out, leaving you a clean slate with which to create new NTFS permissions for the object. Any permissions that were specifically set for this resource beforehand remain. If you copy the settings, all the settings remain the same, but the gray goes away from the checkboxes, indicating that these settings are now independent settings for this folder or file only.

Special access permissions

But wait—there's more. In addition to the normal NTFS permissions, there are 14 “special access” permissions. These let you fine-tune the permissions granted for a particular object. These are not actually separate permissions from the standard ones, but rather refinements of them. For example, the standard Read permission actually involves four separate permissions rolled into one. The special permissions break them down into four separate settings: Read Data, Read Attributes, Read Permissions, and Read Extended Attributes.

By default, the special access permissions are set according to the standard permission settings you have specified, but you can change them as desired.

To view the special permission settings, click the Advanced button on the Security tab. This opens the Access Control Settings For Data dialog box, shown in **Figure B**.

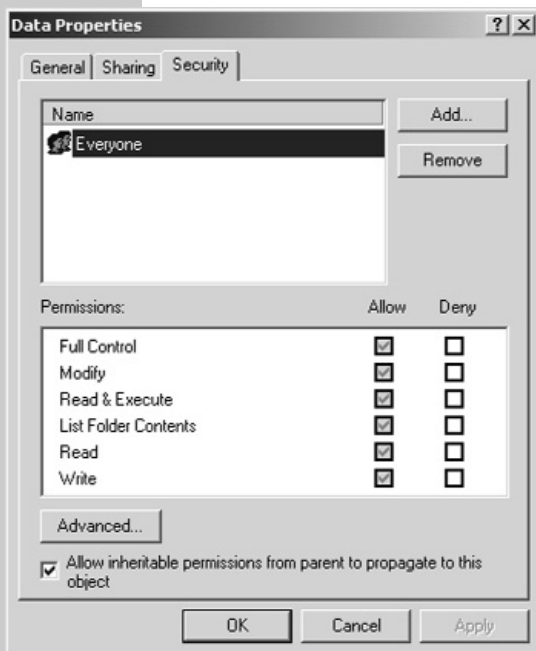
From here, double-click one of the listed users or groups to display the settings for the 14 extra permissions. **Figure C** shows the Permission Entry For Data dialog box that opens.

Most of these special permissions are useful only in odd circumstances. For example, suppose you have granted a group Modify access to a particular folder, but you want to make it impossible for them to delete a certain file in that folder. You could set one of the special access permissions—Delete—to Deny for that file.

Ownership

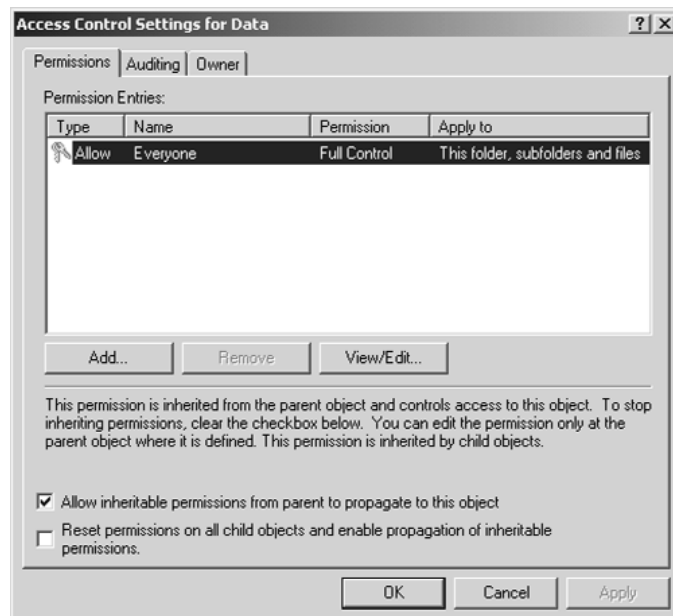
There are two special access permissions you might use more frequently: Change Permissions and Take Ownership. Change Permissions is a permission that normally comes only with Full Access, but you can specifically grant

Figure A



Set NTFS permissions on the Security tab on the Data Properties sheet.

Figure B



Control access for a resource more precisely from the Access Control Settings For Data dialog box.

it for a resource here. Take Ownership allows a user to transfer the ownership of the file or folder to himself or herself. There can be only one “owner” for a file or folder at a time, and that user is the only member of the CREATOR OWNER group for that object. You can assign certain rights to that group, just as you can assign permissions to any other group. The Take Ownership permission enables someone to usurp the title of Owner from another for that resource.

Note that having *permission* to take ownership of a resource does not automatically *take* the ownership. If a user has the permission to take ownership, the Owner tab appears in the Access Control Settings dialog box for the resource. Click the Owner tab and then choose yourself on the list of users. (You cannot choose anyone else; you must choose the user name with which you are logged on.) If you also want to take ownership of all subordinate folders and files, mark the Replace Owner On Subcontainers And Objects check box.

What happens to permissions when you move or copy?

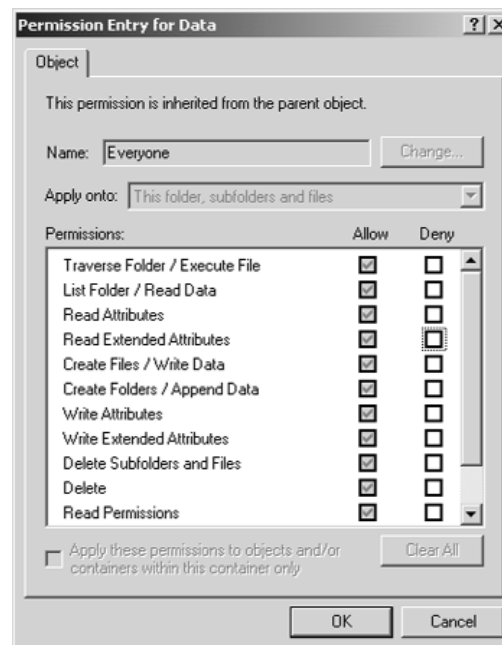
When you copy a folder that has specifically been shared (rather than just inheriting sharing from its parent), the original remains shared, but the copy is reset to Not Shared. However, if you copy the folder to a drive or folder that is shared, it will inherit the sharing setting of its new parent location. The same goes for moving a folder. Any specific sharing permissions it has are removed, but it is free to inherit sharing from the new location.

When you copy or move a file or folder from an NTFS drive to a FAT or FAT32 drive, all NTFS permission settings are removed, leaving it wide-open for anyone to access.

When you copy to another NTFS drive or within the same drive, any old NTFS permissions assigned specifically to the original are stripped away, and it inherits NTFS permissions from the new location. In order to copy, you must have Write permission for the destination. The user doing the copying becomes the CREATOR OWNER of the copy.

When you move a file or folder to another NTFS drive, the permissions work just like

Figure C



You can set more specific permissions here than are possible with the normal NTFS permissions.

HERE ARE SOME MORE TIPS FOR USING NTFS PERMISSIONS

- Try to assign NTFS permissions to folders rather than individual files, and make sure that the files are set to inherit their permission from the folder. (That’s the default setting, so you don’t have to check every single file.)
- Create folders according to access requirements—for example, a folder for files that Marketing needs, another for files that Engineering needs, and so on—and assign NTFS permissions to those folders for the people who need them.
- To prevent users from accidentally deleting important applications or data, remove the Everyone permission and assign the Read & Execute permission to the Users group and the Administrators group for the folder.
- As with sharing permissions, give users only the access level that they require. In most cases, Full Control should reside only with the CREATOR OWNER group.
- Don’t use Deny except when it is necessary, because it can create administrative headaches later.

copying. Any old permissions are removed, and the file or folder inherits permissions from the new location. You must have Modify

permission for the file or folder being moved and Write permission for the destination drive or folder. The user doing the moving becomes the CREATOR OWNER of the file.

When you move a file or folder to a different location on the same NTFS drive, however, permissions work a little differently. The moved file or folder does inherit permissions from the new location, but if there were any permissions set specifically for that object, they are retained and they override the new inheritances. You must have Modify permis-

sion for the file or folder being moved and Write permission for the destination drive or folder. The CREATOR OWNER does not change.

Conclusion

In this article, you learned to create folder and file permissions for groups and individuals using the NTFS file system. You learned how NTFS permissions are inherited and what happens when you move or copy folders and files. ~

Combining sharing and NTFS permissions in Windows 2000

Feb 6, 2001

By Faith Wempen, A+, MOUS 2000 Master

In this article, I cover the tricky subject of what happens when you combine permissions. After reading this and the preceding articles, you should be able to set up and troubleshoot permissions on your network and clients more quickly.

Rules for combining permissions

Understanding how permissions interact is not difficult, if you stick with these rules.

Same permission type (either sharing or NTFS)

When working within a certain permission type (sharing or NTFS), permissions are cumulative. The most lenient setting wins for a particular user or group. Deny always overrides Allow and negates any permission with which it conflicts.

Mixing sharing and NTFS permissions

When there's a difference between the sharing permission and the NTFS permission, the most restrictive setting wins.

Permissions across groups

Permissions are not cumulative across groups; each group's permission is calculated separately. For example, if a user is a member of Group A that has Full Control sharing permission but no NTFS permission for an object and of Group B that has Full Control NTFS permission but no sharing permission for the object, that user has no permission for the object.

Examples

Let's look at some examples. Let's say that on Tim's PC, there is a folder called FOLDER-A containing a file called PRIVATE.DOC. Tim has shared FOLDER-A with the Marketing group with Change permission and with the Everyone group with Read permission. In the NTFS permissions for the folder, he has allowed for the Marketing group to have only Read access. He has removed the default permissions to the folder for the Everyone group.

If Sarah from Marketing accesses PRIVATE.DOC, will she be able to make changes

Table A

	Sharing permission	NTFS permission	Net permission
Marketing group	Change	Read	Read
Everyone group	Read	None	None
Cumulative permission			Read

Table B

	Sharing permission	NTFS permission	Net permission
Marketing group	Change	Read	Read
Managers group	None	Modify	None
Everyone group	Read	None	None
Cumulative permission			Read

Table C

	Sharing permission	NTFS permission	Net permission
Marketing group	Change	Modify	Change/Modify
Managers group	None	Modify	None
Everyone group	Read	None	None
Cumulative permission			Change/Modify

to it? The Marketing group has Change (for Sharing) and Read (for NTFS), with a net result of Read. The Everyone group has Read (for Sharing) and None (for NTFS), with a net result of None. So Sarah's permissions are the least restrictive of Read and None—in other words, Read. So no, she cannot make changes (see **Table A**).

Now, suppose Tim adds another group to his list of NTFS permissions: Managers. He gives the Managers group Modify access to FOLDER-A. If Sarah is a member of the Managers group, will she now be able to make changes to PRIVATE.DOC? The answer is still no, because even though permissions are cumulative within a type, they are calculated as a whole on each group. As you can see above, the new Managers group has no net permission to the folder because it has no Sharing permission, so it doesn't help Sarah to be able to modify the file (see **Table B**).

HINT

Permission changes don't take effect until the end user logs off and back on. After Tim changes the permissions, Sarah must log off and back on again or close the network connection to Tim's PC and reopen it in order for his permission changes to take effect on Sarah's end.

If Tim wanted to make sure Sarah had the ability to modify the file, he could:

- Give the Marketing group Modify (or better) permission under NTFS permissions.
- Give the Managers group Change permission under sharing permissions.

Tim takes the first option and changes the Marketing group's NTFS permission to Modify. Now the chart looks like **Table C**.

Table D

	Sharing permission	NTFS permission	Net permission
Marketing group	Change	Modify	Change/Modify
Managers group	None	Deny Write	Deny Write
Everyone group	Read	None	None
Cumulative permission			Deny Write

Sharing and NTFS permissions use two different terms, Change and Modify, but both allow Sarah to make edits to the file.

Now, suppose Tim uses the NTFS special permissions to deny the Managers group the Write permission. Will Sarah be able to edit the file? No, because the Deny option settings override any Allow settings. Even though the Marketing group still has the rights to edit the file, Sarah is also a member of the Managers group which is specifically denied access (see **Table D**).


If Tim wanted Sarah but nobody else from the Managers group to be able to change the file, he could either remove Sarah from that group or create a separate group containing everyone from Managers except Sarah and deny that group the Write access instead of denying the Managers group.

Practice

The best way to get more confident in your understanding of permissions is to play around with them. Try re-creating the preceding scenario on two client PCs on your network and then experimenting with more “what if” scenarios. For example, what if:

- ▶ Tim turns off Deny Write for Managers and simply deselects the Allow check box for the Managers group? Can Sarah then edit the file?
- ▶ Sarah then tries to delete the file PRIVATE.DOC? Can she do it with her current permissions?
- ▶ Tim removes all permissions from the folder? Can he still read and modify the file himself?
- ▶ Sarah creates a subfolder within FOLDER-A on Tim’s PC? Can Tim delete it?

Conclusion

In this article, you learned what the rules are when different sets of permissions interact. You also gained some practice in determining net permissions when NTFS and sharing permissions conflict for a user in multiple groups. You now have my permission to set up your network and client machines for the most robust security obtainable in a Windows environment. 

Establish the correct file-sharing permissions in Windows XP

Apr 8, 2002

By TechRepublic Staff

With the NT file system (NTFS) in Windows XP, you can set file permissions at the local PC level in addition to the file-sharing permissions of the network environment. Along with this additional functionality comes complexity and the potential for all kinds of admin headaches. One harried manager wants to know why he can't access data on a colleague's PC that he needs to assemble a presentation; another can't figure out why the mailroom intern was able to browse the files he thought he had secured. More options mean more chances for confusion and user error, and if you don't have a thorough understanding of the various permissions and their relationships, it can be nearly impossible to sort out a permission problem and find a solution.

We'll review the file and folder permissions in Windows XP. Once you understand Windows XP permissions and how they interact, you'll be able to troubleshoot permission issues that occur on your network more quickly.

Watch file-sharing and NTFS permission interactions

In any Windows network environment (peer-to-peer or server-based), you can set sharing permissions for drives and folders. By default, when you set up a PC on a network, no drives or folders on that PC are shared. The local user of that PC can then choose to share entire drives or individual folders on a drive. This type of security is not really that secure, however, because it affects only network access. Local access (that is, someone sitting down at the PC and logging on) is wide open.

For drives formatted with NTFS, you can also set NTFS permissions. These can affect drives and folders as well as individual files. NTFS permissions affect local users as well as network users and are based on the permission granted to individual user logons, regardless of where they're connecting. You also have a

much wider variety of permissions to choose from with NTFS permissions, so you can more precisely control the rights being granted.

When file sharing permissions and NTFS permissions conflict, the most restrictive of the two wins. For example, if someone has full access to a certain file from NTFS permissions but has no sharing permissions to the folder in which it resides, he or she cannot access the file from the network. He or she can, however, physically sit down at the local PC containing the file, log in, and access it, because sharing permissions do not affect local access.

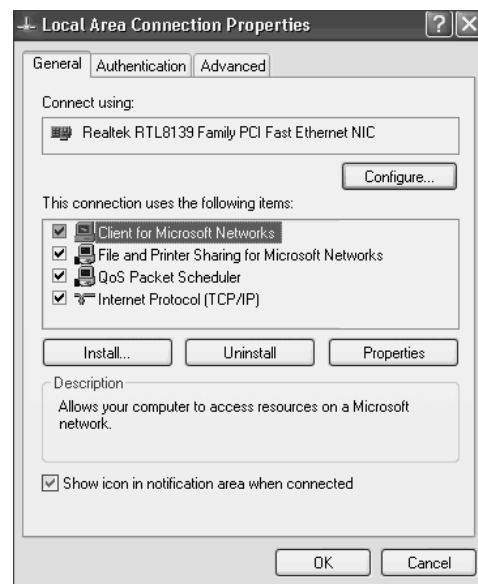
Working with shared folders

Shared folders provide remote access to the files on a PC. Folder sharing is available on drives using all types of partitions: FAT, FAT32, or NTFS. To share any folders (or any printers, for that matter) on a Windows XP PC, File And Printer Sharing For Microsoft Networks must be installed as a networking component. To check for it, right-click the Local Area Connection icon in the Windows XP taskbar and choose Status. From the Local Area Connection Status dialog box, select the Properties button to see the listing shown in **Figure A**. If File And Printer Sharing For Microsoft Networks doesn't appear on the list, add it by clicking the Install button and choosing it from the Services category.

After File And Printer Sharing For Microsoft Networks is in place, you can share individual drives and folders by right-clicking a drive or folder and choosing Sharing And Security. When you do, the Sharing tab of the Properties dialog box will open.

Sharing is slightly different for drives than for files. With a drive, you might see a default share already set up. These have a dollar sign (\$) following the share name, as shown in **Figure B**. Such shares are for administrative use only; ordinary users won't be able to see or

Figure A




File And Printer Sharing For Microsoft Networks must be installed to share folders over a network.

Figure B



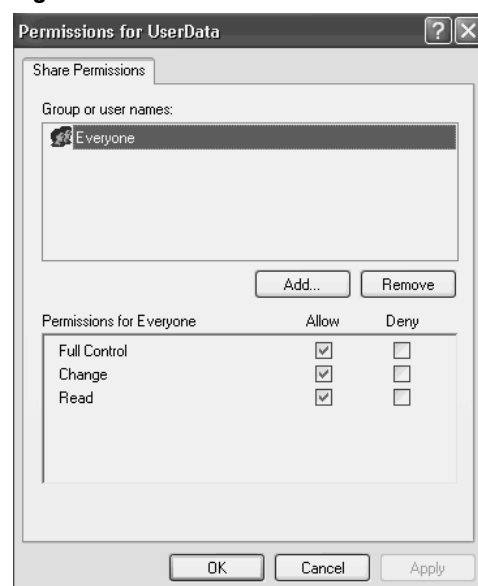
C\$ is the default administrative share for this drive; it doesn't count as a user-to-user share.

Figure C



Create a new share to allow other users to access the drive.

Figure D



Limit permission to the folder or drive, if desired.

browse a drive shared in this way on the network. Consequently, if you want to share an entire drive like this on your network, you must create an additional share for it.

To create a new share for a drive, click the New Share button and then fill in the Share Name, any comment you want to make, and a user limit for concurrent usage (if desired). While you're in the New Share dialog box (see **Figure C**), you can click the Permissions button to specify who will have access to the shared drive, or you can save that for later.

For a folder, the process is more straightforward because there are no default administrative shares. By default, a folder is set to Do Not Share This Folder. To share it, right-click the folder and select Sharing And Security from the context menu. Choose the Share This Folder button and then enter a share name, comment, and user limit.

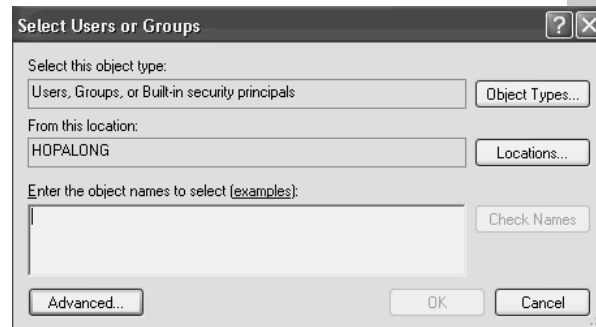
Regardless of whether you're sharing a folder or a drive, you can configure permissions the same way: Display the Sharing tab and click the Permissions button. A Permissions dialog box will appear, as shown in **Figure D**. By default, all permissions are granted to everyone.

If you plan to use NTFS permissions in conjunction with sharing permissions, you might want to leave the sharing permissions

TIP: DON'T DENY

The Deny option should be used sparingly because it overrides any more lenient permissions. For example, if you set Read access for a folder to Deny and the drive on which the folder resides allows Full Control, everything on that drive will have Full Control access except for that folder, which will have no access at all.

Figure E



Specify other users or groups to receive permissions.

TIPS FOR USING SHARING PERMISSIONS EFFECTIVELY

- ▶ Grant only the permissions that a group or user needs; disallow all others. In most cases, Change permission is all a user needs for a drive or folder. Change enables users to run programs, edit files, and so on.
- ▶ Don't allow Full Control for a drive to the Everyone group. If certain users must have complete control of a drive, assign Full Control to a particular group or create a group for that purpose.
- ▶ Don't use the Deny option unless you have a specific reason to do so. It's easy to forget that you've used the Deny option and spend fruitless hours troubleshooting a file access issue because of it.
- ▶ Assign sharing permissions to groups, rather than individuals, to minimize administrative work.
- ▶ Use descriptive share names to help users locate the shared drives or folders they want.
- ▶ Group the folders that need to have the same sharing permissions assigned together in a single folder and then assign the permissions to the parent folder.

set at the default “free-for-all” settings and rely on the NTFS permissions to lock down certain sensitive items. However, if you aren't going to use NTFS permissions or if you can't because the drive is FAT or FAT32, you can restrict access at the sharing level.

Note in Figure D the three types of sharing permissions:

- ▶ **Read:** Users can display the contents of the folder, open files, display attributes, and run programs.
- ▶ **Change:** Users have all the rights of Read, plus the ability to create new folders and files within the shared folder or drive, open and change files, change file attributes, and delete folders and files.
- ▶ **Full Control:** Users have all of the rights of Change, plus the ability to take ownership of files and change file permissions.

Everything within a shared drive or folder inherits its sharing permissions. For example, if

a shared drive has 10 folders, all of those folders have the same sharing permissions as the drive, unless they are set otherwise. Permissions are cumulative, which means that, in the event of a conflict between a specific folder's permissions and those it has inherited from the drive (or parent folder), the most lenient wins. For example, if you allow Read access on a folder and don't allow Change or Full Control on that folder, but the drive itself allows Full Control, that folder will also have Full Control access permitted.


For each setting (Read, Change, and Full Control), you can choose the option to Allow or Deny. The default is set to Allow. If you don't want to allow a particular permission, you simply deselect the Allow check box. “Disallowing” something (that is, turning off Allow permissions for it) takes away that right but enables the folder to inherit permissions from the parent folder or drive.

When you share a folder or drive, only one group has permissions assigned by default: the Everyone group. That means all users will have the same permission rights to the object, regardless of any group affiliation. You can delete the Everyone group from the list and/or add other groups or individuals to the permissions list. You might, for example, delete the Everyone group from the list entirely or leave it there and set it to allow Read permission only and then add the Administrators group to the list and grant that group Full Control.

To add a group or user to the permissions list for an object, start from the Permissions dialog box (Figure D), click the Add button, type the user or group you want in the Select Users Or Groups dialog box (Figure E), and click the OK button. If you don't know the exact name of the group or user, click on the

Advanced button and select Find Now to perform a search on the available choices. When you're finished, click OK to return to the Permissions dialog box. The users and groups you chose will appear on the Permissions list, ready to have their permission levels set.

Get file permissions right the first time

The proper sharing of files on a network is of extreme importance to you, the network administrator. Without a thorough understanding of how Microsoft configures file sharing, you'll find your users making daily demands of your time to fix file access problems. The next two articles will specifically address NTFS permissions in Windows XP and using the two types of permissions effectively. 

Effectively set and troubleshoot NTFS permissions in Windows XP

Apr 10, 2002

By TechRepublic Staff

Setting folder and file permissions gives you some network security, but it doesn't secure your PC desktop. When you use the NT file system (NTFS) in Windows XP, however, you can set file permissions at the local PC level. That means that a user sitting down at a PC—not just a user accessing the resource across a network—is bound by NTFS permissions.

NTFS permissions, which can be set only on drives partitioned with NTFS, can be assigned to drives and folders, just like sharing permissions, but they also can be assigned to individual files. Unlike sharing permissions, in which the default setting for a resource is Not Shared, NTFS permissions are set to allow access by default.

In this article, we'll cover the details of NTFS permissions in Windows XP. With an understanding of how NTFS permissions work, you'll be able to troubleshoot permission issues more quickly as they occur on your network and clients.

Folder and drive permissions

NTFS offers many more types of permission than the simple Read, Change, and Full Control of sharing permissions. For folders and drives, you can assign these permissions:

- **List Folder Contents:** View a folder's contents
- **Read:** View a folder's contents, open files, and view file and folder attributes

- ▶ **Read & Execute:** Same as Read, plus the ability to move through folders to reach other folders, even if no permission is granted for those folders
- ▶ **Write:** Same as Read, plus the ability to create and edit subfolders and change attributes
- ▶ **Modify:** Combination of Read & Execute and Write, plus permission to delete the folder
- ▶ **Full Control:** Same as Modify, plus the ability to change permissions, take ownership, and delete subfolders and files
- ▶ **Special Permissions:** Allows you to customize permissions on folders by selecting the individual components of the standard sets of permissions

File-level permissions

The list of permissions for individual files is the same, except for the List Folder Contents permission. For files, you can assign these permissions:

- ▶ **Read:** Open the file and view its attributes, ownership, and permissions
- ▶ **Read & Execute:** Same as Read, plus the ability to run applications
- ▶ **Write:** Same as Read, plus the ability to change file content and attributes
- ▶ **Modify:** Same as Write and Read & Execute combined, plus the ability to delete the file
- ▶ **Full Control:** Same as Modify, plus the ability to change permissions and take ownership
- ▶ **Special Permissions:** Allows you to customize permissions on files by selecting the individual components of the standard sets of permissions

Just like sharing permissions, NTFS permissions can be set to Allow with the Allow check box. Permissions are cumulative and can be inherited from parent folders or drives. NTFS permissions can also be set to Deny, but you should use Deny sparingly because it overrides more lenient permissions. For example, if you set Read access for a folder to Deny and the drive on which the folder resides allows Full

Control, everything on that drive will have Full Control access except for that folder, which will have no access at all.

To set NTFS permissions, use the Security tab on the Properties page for a drive, folder, or file (see **Figure A**). The controls will seem familiar; they're almost the same as the ones for setting sharing permissions.

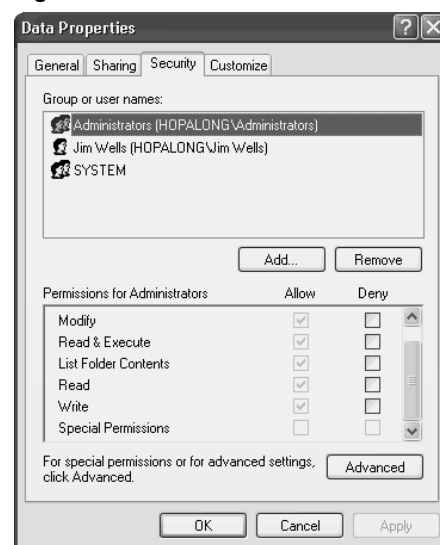
Special access permissions

In addition to the normal NTFS permissions, you can use 14 “special access” permissions. These let you fine-tune the permissions granted for a particular object. They're not actually separate permissions from the standard ones but refinements of them. For example, the standard Read permission actually involves four separate permissions rolled into one. The special permissions are the four separate settings: Read Data, Read Attributes, Read Permissions, and Read Extended Attributes. By default, the special access permissions are set according to the standard permission settings you have specified, but you can change them as desired.

To view the special permission settings, click the Advanced button on the Security tab to open the Advanced Security Settings For Data dialog box, as shown in **Figure B**.

From here, double-click one of the listed users or groups to display the settings for the 14

Figure A



Set NTFS permissions on the Security tab on the Data folder's Properties box.

extra permissions. **Figure C** shows the Permission Entry For Data dialog box that will open.

Most of these special permissions are useful only in odd circumstances. For example, suppose you have granted a group Modify access to a particular folder, but you want to make it impossible for them to delete a certain file in that folder. You could set one of the special

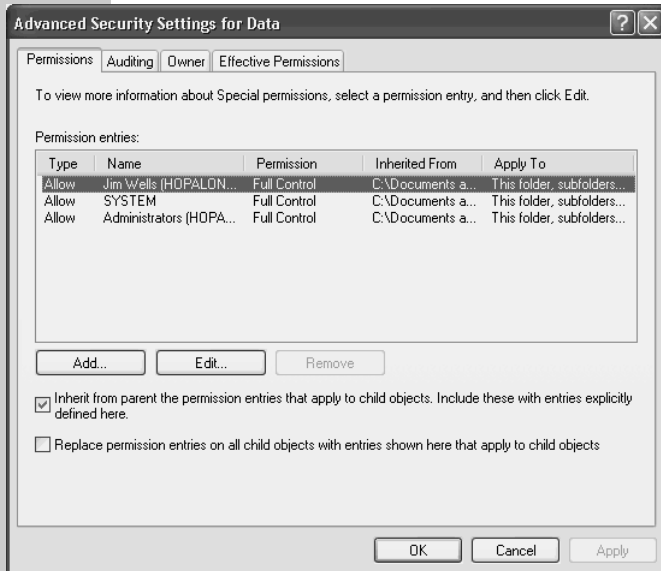
access permissions—Delete—to Deny for that file.

Inheriting permissions

Notice the first check box at the bottom of Figure B. When this option is turned on, the folder or file will inherit the permissions of the parent object (that is, the drive or folder in which it resides). The grayed-out check boxes in Figure C indicate that those permissions are inherited rather than specific to this folder.

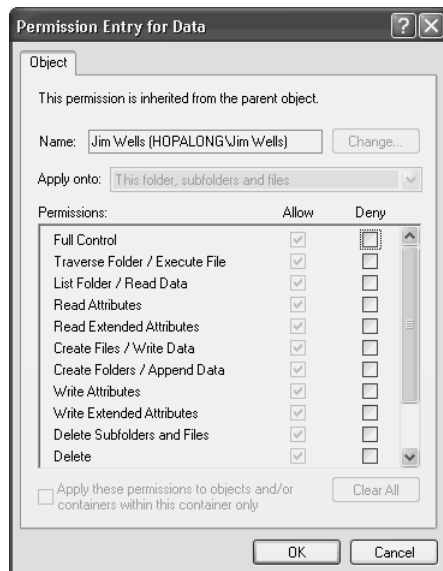
If you deselect the Inherit From Parent The Permission Entries check box, a dialog box will ask what you want to do about those inherited settings. (You won't see this on drives, because they have nothing to inherit from, being at the top level already.) You can choose to copy them or to remove them. If you remove them, all permissions and all users that were inherited are stripped out, leaving you a clean slate with which to create new NTFS permissions for the object. Any permissions that were specifically set for this resource beforehand remain. If you copy the settings, all the settings remain the same, but the check boxes become active, indicating that these settings are now independent settings for this folder or file only.

Figure B



Control access for a resource more precisely from the Advanced Security Settings For Data dialog box.

Figure C



You can set more specific permissions here than are possible with the normal NTFS permissions.

Ownership

You might use two special access permissions more frequently: Change Permissions and Take Ownership. You can find the Change Permissions feature on the Effective Permissions tab of the Advanced Security Settings For Data dialog box. Change Permissions is a permission that normally comes only with Full Access, but you can specifically grant it for a resource here.

Located on the Owner tab of the Advanced Security Settings For Data dialog box, Take Ownership allows a user to transfer the ownership of the file or folder to himself or herself. There can be only one “owner” or a file or folder at a time, and that user is the only member of the CREATOR OWNER group for that object. You can assign certain rights to that group, just as you can assign permissions to any other group. The Take Ownership permission enables someone to

MORE TIPS FOR USING NTFS PERMISSIONS

- ▶ Try to assign NTFS permissions to folders rather than individual files and make sure that the files are set to inherit their permission from the folder. (That's the default setting, so you don't have to check every single file.)
- ▶ Create folders according to access requirements—for example, a folder for files that Marketing needs, another for files that Engineering needs, and so on—and assign NTFS permissions to those folders for the users who need them.
- ▶ To prevent users from accidentally deleting important applications or data, remove the Everyone permission and assign the Read & Execute permission to the Users group and the Administrators group for the folder.
- ▶ As with sharing permissions, give users only the access level that they require. In most cases, Full Control should reside only with the CREATOR OWNER group.
- ▶ Don't use Deny except when it is necessary, because it can create administrative headaches later.

usurp the title of Owner from another for that resource.

Note that having *permission* to take ownership of a resource does not automatically *take* the ownership. If a user has the permission to take ownership, click the Owner tab and then choose yourself on the list of users. (You cannot choose anyone else; you must choose the user name with which you are logged on.) If you also want to take ownership of all subordinate folders and files, select the Replace Owner On Subcontainers And Objects check box.

What happens to permissions when you move or copy?

When you copy a folder that has specifically been shared (rather than just inheriting sharing from its parent), the original remains shared, but the copy is reset to Not Shared. However, if you copy the folder to a drive or folder that is shared, it will inherit the sharing setting of its new parent location. The same goes for moving a folder. Any specific sharing permissions it has are removed, but it's free to inherit sharing from the new location.

When you copy or move a file or folder from an NTFS drive to a FAT or FAT32 drive, all NTFS permission settings are removed, leaving it wide-open for anyone to access.

When you copy to another NTFS drive, or within the same drive, any old NTFS permissions assigned specifically to the original are stripped away, and it inherits NTFS permissions from the new location. To copy, you must have Write permission for the destination. The user doing the copying becomes the CREATOR OWNER of the copy.

When you move a file or folder to another NTFS drive, the permissions work just as they do when you copy them: Any old permissions are removed, and the file or folder inherits permissions from the new location. You must have Modify permission for the file or folder being moved and Write permission for the destination drive or folder. The user doing the moving becomes the CREATOR OWNER of the file.

When you move a file or folder to a different location on the same NTFS drive, however, permissions work a little differently. The moved file or folder does inherit permissions from the new location, but if any permissions were set specifically for that object, they're retained and they override the new inheritances. You must have Modify permission for the file or folder being moved and Write permission for the destination drive or folder. The CREATOR OWNER doesn't change.

NTFS means more permissions options

Windows XP NTFS permissions features allow greater control for you and more configuration schemes for your users. In this article,

you learned to create folder and file permissions for groups and individuals using the NTFS file system. You also learned how NTFS permissions are inherited and what happens when you move or copy folders and files. ~

Combining sharing and NTFS permissions in Windows XP

Apr 11, 2002

By TechRepublic Staff

In this article, we cover the tricky subject of what happens when you combine NTFS and file-sharing permissions in Windows XP. After reading this article, you'll be able to set up and troubleshoot permissions on your network and client more quickly.

Rules for combining permissions

Understanding how permissions interact isn't difficult if you stick with these rules:

- ▶ When working within a certain permission type (sharing or NTFS), permissions are cumulative. The most lenient setting wins for a particular user or group. Deny always overrides Allow and negates any permission with which it conflicts.
- ▶ When there's a difference between the sharing permission and the NTFS permission, the most restrictive setting wins.
- ▶ Permissions are not cumulative across groups; each group's permission is calculated separately. For example, if a user is a

member of Group A, which has Full Control sharing permission but no NTFS permission for an object, and also of Group B, which has Full Control NTFS permission but no sharing permission for the object, that user has no permission for the object.

Examples

Let's look at some examples. Say that on Tim's PC is a folder, FOLDER-A, containing a file, PRIVATE.DOC. Tim has shared FOLDER-A with the Marketing group with Change permission and with the Everyone group with Read permission. In the NTFS permissions for the folder, he has allowed for the Marketing group to have only Read access. He has removed the default permissions to the folder for the Everyone group. If Sarah from Marketing accesses PRIVATE.DOC, will she be able to make changes to it? The Marketing group has Change (for sharing) and Read (for NTFS), with a net result of Read. The Everyone group has Read (for sharing) and None

Table A

	Sharing permission	NTFS permission	Net permission
Marketing group	Change	Read	Read
Everyone group	Read	None	None
Cumulative permission			Read

Table B

	Sharing permission	NTFS permission	Net permission
Marketing group	Change	Read	Read
Managers group	None	Modify	None
Everyone group	Read	None	None
Cumulative permission			Read

(for NTFS), with a net result of None. So Sarah's permissions are the least restrictive of Read and None—in other words, Read. So no, she cannot make changes (see **Table A**).

Now, suppose Tim adds another group to his list of NTFS permissions: Managers. He gives the Managers group Modify access to FOLDER-A. If Sarah is a member of the Managers group, will she now be able to make changes to PRIVATE.DOC? The answer is still no, because even though permissions are cumulative within a type, they're calculated as a whole on each group. As you can see below, the new Managers group has no net permission to the folder because it has no sharing permission, so it doesn't enable Sarah to modify the file (see **Table B**).

If Tim wanted to make sure Sarah had the ability to modify the file, he could:

HINT

Permission changes don't take effect until the end user logs off and logs back on. After Tim changes the permissions, Sarah must log off and back on again or close the network connection to Tim's PC and reopen it in order for his permission changes to take effect on Sarah's end.

- ▶ Give the Marketing group Modify (or better) permission under NTFS permissions.
- ▶ Give the Managers group Change permission under sharing permissions.

Let's say Tim takes the first option and changes the Marketing group's NTFS permission to Modify. Now the chart looks like **Table C**.

Table C

	Sharing permission	NTFS permission	Net permission
Marketing group	Change	Modify	Change/Modify
Managers group	None	Modify	None
Everyone group	Read	None	None
Cumulative permission			Change/Modify

Table D

	Sharing permission	NTFS permission	Net permission
Marketing group	Change	Modify	Change/Modify
Managers group	None	Deny Write	Deny Write
Everyone group	Read	None	None
Cumulative permission			Deny Write

NOTE

Sharing and NTFS permissions use two different terms, **Change and Modify**, but both allow Sarah to make edits to the file.

Now, suppose Tim uses the NTFS special permissions to deny the Managers group the Write permission. Will Sarah be able to edit the file? No, because the Deny option settings override any Allow settings. Even though the Marketing group still has the right to edit the file, Sarah is also a member of the Managers group, which is specifically denied access (see **Table D**).


If Tim wanted Sarah, but nobody else from the Managers group, to be able to change the file, he could either remove Sarah from that group or create a separate group containing everyone from Managers except Sarah and deny that group Write access instead of denying the Managers group.

Practice

The best way to get more confident in your understanding of permissions is to play around with them. Try re-creating the preceding scenario on two client PCs on your network and then experimenting with more “what if” scenarios. For example, what if:

- ▶ Tim turns off Deny Write for Managers and simply deselects the Allow check box for the Managers group? Can Sarah then edit the file?
- ▶ Sarah then tries to delete the file PRIVATE.DOC? Can she do it with her current permissions?
- ▶ Tim removes all permissions from the folder? Can he still read and modify the file himself?
- ▶ Sarah creates a subfolder within FOLDER-A on Tim’s PC? Can Tim delete it?

You have our permission

You’ve now learned what the rules are when different sets of permissions interact. You also gained some practice in determining net permissions when NTFS and sharing permissions conflict for a user in multiple groups. You now have our permission to set up your network and client machines for the most robust security obtainable in a Windows environment. 

Notes

Notes

Wireless Security

Keep up with public wireless dangers and Wi-Fi security standards	95
Design a secure wireless LAN.....	97
Think security when setting up an 802.11b wireless network	100
How to beef up wireless security	102
Use WEP to improve security on your wireless network	105
Take steps to secure vulnerable WLANs	110
At last, real wireless LAN security	111
WPA wireless security offers multiple advantages over WEP	113
Six tips for implementing closed networking on a wireless network	115
Don't use MAC filtering as your only wireless network security solution	116
Choosing a vendor solution for wireless LAN security with 802.1x and EAP	119
Follow these steps to tighten security on Linksys wireless networks	121
XP client configuration for enhanced security on a Linksys wireless network.....	124

Keep up with public wireless dangers and Wi-Fi security standards

Aug 11, 2003

By John McCormick

Although wireless networking holds great promise for extending and mobilizing the 24/7 connected world we've all become accustomed to, it obviously comes with a wide variety of manageability and security headaches for IT departments. Two of the biggest problems IT administrators currently face are protecting mobile users who are now connecting to public wireless hotspots and keeping well informed about the latest standards and techniques for securing wireless LANs.

The public wireless problem

More and more wireless networks now beckon the unwary road warrior. So it's become vital for administrators to take responsibility for the mobile workers carrying company data out into the connected world of airports, high dollar coffee shops, hotels, and restaurants and taverns—many of which now allow users to connect their laptops and/or PDAs to the Web using wireless public networks.

If you've never given this a thought before, consider how little your laptop-equipped users are aware of the dangers of logging on to any random network they encounter in their travels. At a bare minimum, you need to educate them about the threat these open networks pose. You may also need to scrub their systems of any critical unencrypted corporate data they are carrying around.

Just as companies are coming to realize how dangerous unfiltered access to the Internet is in the office, IT professionals as well as users must start viewing public wireless networks as a wilderness where many systems could become easy prey for attackers. After all, why should a hacker go to all the trouble of breaking into a corporate network when an open wireless network provides easy access to a corporate system? From there, an attacker can, for example, plant a Trojan or raid corporate data stored locally on the system.

A well-configured firewall is essential for any laptop that has wireless capabilities—

regardless of whether the person carrying it has any confidential information—because, at a minimum, they may pick up a Trojan, a virus, or other malicious software and later transfer it to the company network.

Keep up with WLAN security

Securing your own wireless network can be a much bigger challenge than guarding your mobile users, and this is due both to weak security offerings and a confusion of standards in the wireless field. In fact, most wireless vendors ship their offerings with encryption turned off and/or with very weak security settings as part of the default configuration.

Even with encryption turned on, a Wi-Fi network is inherently insecure because the encryption used is weak. Forcing your users to use encryption locally will at least prevent the average script kiddie—who just got a laptop as a birthday present—from penetrating your system by doing little more than walking past your office building. The effort to encrypt your WLAN may also provide a good legal, if not technical, defense against serious hackers taking over your network for illegal purposes.

Although configuring an open wireless LAN has become so simple that virtually anyone can do it, securing one is a major challenge worthy of the time and talents of a top security expert.

In the beginning, 802.11b relied primarily on MAC address filtering for access control. If you had an allowed MAC address, you could connect to the wireless access point. The only problem was spoofing. Your wireless device was continually broadcasting its address and any attacker could intercept it and spoof the MAC address to match the allowed address.

Data was secured using Wired Equivalent Privacy (WEP). But WEP generally uses a 40-bit encryption key (sometimes 64-bit) and only a 24-bit initialization vector (IV), which makes it extremely vulnerable. The 128-bit WEP2 is available on some systems. A major problem

A TESTAMENT TO THE WIRELESS PROBLEM

I decided to write this article after I spoke at the Summercon hacker convention in Pittsburgh recently. There were probably 30 open networks within a single square mile around the conference site, and other cities have similar WLAN-rich areas around universities and high-tech businesses. I saw people logging on to three and four wireless networks from PDAs right in the hotel lobby, and only one of the networks was owned by the hotel.

Everyone from the overt FBI agent to a former NSA staff member to the average hacker was logging on to wireless networks, and I bet even in that elite group, no more than half realized that merely by connecting to an open network they were potentially opening up their computers to anyone else on the same wireless network.

Even worse, only a few of those networks were intended for general public use. Most were private networks with so little security that anyone could log on, almost by accident.

with WEP is the 24-bit IV, which is so small that many networks will reuse the same IV multiple times in a single day. In fact, it is so insecure that there are free hacker tools available on the Internet to crack a busy WEP network in a few hours.

Adding IPSec can be a major improvement for security, but most wireless networks are already plagued with quality of service (QoS) issues, so using sophisticated encryption schemes across the network is usually not an acceptable solution unless you upgrade the hardware on the entire network.

Several wireless vendors have quickly moved to secure their market share by improving the authentication side to offer better security for their products. Cisco and Microsoft have pushed RSA's Protracted Extensible Authentication Protocol (PEAP) to authenticate users through a secure tunnel. Cisco also has another security protocol, termed Lightweight EAP (LEAP), which is simple to implement (on Cisco equipment), although the passwords may be vulnerable to dictionary attacks. Both of these are based on the IEEE 802.1X framework and are improvements over WEP, allowing authentication without having a certificate on the client.

But PEAP isn't as useful as it could be because Cisco's version isn't the same as Microsoft's and—surprise—the two aren't compatible. The EAP-TLS protocol used in Windows XP's 802.1X client is strong but requires both server and client certificates.


Another EAP-based protocol, Tunnelled Transport Layer Security (TTLS), developed

by Funk Software, is nearly identical to PEAP—but the key word is “nearly.” EAP-TTLS offers strong security and easy configuration, requiring only server-side certificates.

The new Wi-Fi Protected Access (WPA) is also being pushed by Microsoft, Cisco, and members of the Wi-Fi Alliance. You can download a WPA upgrade for Windows XP from Microsoft (<http://www.microsoft.com/whdc/hwdev/tech/network/802x/WPA.msp>).

None of these EAP-based authentication systems fully address the data security problem posed by the continued reliance on WEP, which is why many organizations have turned to using VPNs to encrypt all communications sent over a wireless link. The problem with that is that it adds another layer of latency and complexity to the WLAN. In short, it simply shouldn't have to be that difficult to make a secure WLAN connection.

Final word

This only skims the surface of the protocol wars raging in the wireless world at this time. In the near term, if you're adding (or moving entirely to) a wireless network, you'll be well advised to stick with a single vendor throughout if you hope to secure your wireless networks. Otherwise, you need to choose technologies compatible with some third-party vendor and rely on that company to keep your system working. Even if you get everything working properly, you should still take a long, hard look at what information you place on that network. 

Design a secure wireless LAN

Sep 26, 2002

By Scott Lowe, MCSE

Wireless LANs allow both legitimate users and hackers to access your network quickly and easily. By securing your wireless LAN, you can avoid opening your network doors to hackers. In this article, I'll show you what you need to do. A wireless network can allow you and your users to work in a significantly more flexible and convenient manner while still reducing infrastructure costs, but it can also create a number of major security issues that must be addressed when the system is installed. The same flexibility that makes wireless so attractive to your company can allow hackers to leave a giant hole in your otherwise secure network. To keep your wireless network free from security breaches, you need to focus on security from inception to implementation.

What are the risks?

First, it's important to understand the security issues that arise with the use of a wireless network. Because a wireless network is accessed via strategically placed antennas, you no longer have specific points of network access like you do with a wired network.

There are a number of security risks associated with the currently widespread 802.11b and 802.11a wireless standards. 802.11b devices operate at up to 11 Mbps while 802.11a devices operate at up to 54 Mbps. Both standards operate using Wired Equivalency Protocol or Wired Equivalent Privacy (WEP), which provides some measure of security for transmission over the airwaves.

There is no one action to take that would secure your wireless LAN. Instead, you'll need to rely on a number of different actions that will offer a multifaceted approach to wireless security.

Unauthorized usage (aka Insertion Attack)

Perhaps one of the biggest problems with improperly secured wireless networks is their ability to be used by anyone within antenna range—even people outside the building. This

is generally not an issue for wired networks, because you know exactly the points through which a potential user could access your network, and you can use security devices such as firewalls to protect against unwanted traffic from outside the network.

The problem of unauthorized users gaining access to unsecured wireless access points is exacerbated by folks who drive around and mark buildings with certain symbols indicating that there is an open wireless network in the area; a practice that has become known as “war chalking” mimicking the old term “war dialing” from the modem days of lore. These symbols have recently caught the attention of the FBI in certain areas, however.

What can you do to make sure that your network does not become identified as a free access point? First, make sure that none of your employees is running a wireless access point that you don't know about. Before you dismiss this as something that you don't need to worry about, ask around. This practice is pretty widespread.

Second, try to position your wireless access point antennas in such a way that communication outside a building in public is minimized. A lot of this is trial-and-error, so be prepared to spend a lot of time finding an optimal location.

Third, you can begin to make use of “authorization lists” based on such information as the MAC address of the wireless NIC. This would require the administrator to keep a list of all of the potential wireless devices that would access the network and to make sure that the wireless access points allow only those devices. Obviously, this creates additional administrative overhead to keep the list up-to-date, but it does help you limit the types of devices that connect to your network. Just keep in mind that that MAC addresses can be spoofed. Anyone with a sniffer would be able to just sit and listen to traffic coming from the wireless access point and eventually get an authorized MAC address that they could then use to gain entry. Therefore, don't just assume

you're secure because you're limiting access based on MAC addresses.

Treat wireless access points as untrusted until you have reason to believe that they are completely secure. You may even want to consider segregating wireless network access on a portion of the network that is separated from the main network by a firewall.

WEP is severely flawed

There are currently three different “standard” security systems in place for wireless networking: WEP, 802.1X, and 802.11i. The most widespread system in place is currently WEP. WEP is the encryption method that is used between the base station and the mobile device to provide a modicum of secure communication. Most WEP-capable devices support either 40 or 128 bit encryption. Although WEP is supposed to secure networks, security professionals have identified extremely dangerous holes in WEP.

WEP uses the RC4 encryption algorithm. This algorithm takes a key and generates a number of pseudo-random keys based on it in order to provide the encryption. Because of the fact that Ethernet is a collision-based networking system, collisions will definitely occur, even with wireless.

Unfortunately, WEP reinitializes the entire data stream after a collision occurs. While someone just walking by with a wireless adapter may be discouraged by the fact that you are using an encrypted data stream, a determined hacker needs only a matter of hours before he or she is able to read enough air traffic to generate the required WEP key to gain access to your network. This applies to both 40- and 128-bit WEP encryption—within similar time frames. This implies that 128-bit WEP encryption is no better than 40-bit which, unfortunately, is the case.

In addition, there are now tools such as AirSnort and WEPCrack that make this job even easier for hackers. AirSnort works by passively listening to traffic. Once it acquires 5-10 million packets, it can guess the encryption password in under a second.

While I recommend that you use WEP to at least prevent less-prepared hackers from gaining access to your network, you

should not count on it as your only source of security.

Slightly newer than WEP, 802.1X is the “second try” for wireless security and has also been proven to have significant security problems, such as being susceptible to session hijacking and man-in-the-middle attacks. Session hijacking involves taking over the session for a client that has already authenticated while man-in-the-middle attacks take advantage of 802.1X's one-way authentication by inserting a node between the wireless client and an access point. While an improvement, 802.1X is not a replacement for WEPs; it simply provides authentication services, not the encryption services that WEP provides.

Currently in the works, the 802.11i standard starts with 802.1X and adds significant features to fix its problems. Most importantly, it adds a key distribution infrastructure that replaces static WEP keys. This will be a huge improvement over WEP. In addition, it is slated to make use of AES (Advanced Encryption Standard) rather than WEP's 40- or 128-bit RC4-based encryption algorithm. For more information on how AES works, check out <http://home.ecn.ab.ca/~jsavard/crypto/co040801.htm>. 802.11i is due by the end of 2002.

So, if WEP is not sufficient, how do you make sure that the traffic that is going out over the air is protected? One way is to use encryption just as you would on a wired network by using such tools as VPNs, SSH, and SCP rather than direct network connections, telnet, and FTP. In fact, making use of a VPN from the wireless client may be an excellent idea in any case because VPNs are a much more well-known element than are wireless networks, and their security issues are much better understood, making them much easier to patch and monitor.

However, you should keep in mind that there are tools that allow wireless hackers to hijack SSH and SSL sessions, thereby invalidating the security that they provide. Often, the only way that users are made aware that this has happened is when the server they are connected to indicates that the host key has changed. If this message is ignored, the hacker has achieved his goal.

Using default SSIDs

A Service Set Identifier (SSID) is a 32-bit character identifier in the header of packets sent over a wireless LAN and acts as a rudimentary identification, password, and authorization mechanism for access to the network. Clients attaching to a base station with SSID enabled must use the same SSID on their clients in order to make use of it.

Out of the box, most vendors' 802.11 gear is useless from a security perspective. All of the default configurations are well known and published all over the Web. For example, Cisco's default SSID is either "tsunami" or "2" with no default telnet password.

If someone just buys the access point, sticks the antenna up in the air, and turns it on without making changes to the defaults, he or she has given a potential hacker access to the network. Therefore, it is important to make sure to change the SSID to something unique and not easily guessed, and to enable passwords for telnet and any other remote administration services. Of course, a hacker listening passively to traffic will eventually be able to get it, but having it enabled could still thwart the attempt.

The SSID is also required for people who need to use the access point, which means that the SSID could be illegally obtained by stealing it from the people who need it for access or by using a stolen wireless device.

Wireless security doesn't have to be an oxymoron

Designing a secure wireless network is a complex task that will result in a lot of work for the administrator that implements it. To keep your wireless network secure, you'll need to endure significant planning and decision-making sessions. To make your planning a little easier, here is a final look at the steps I've covered in this article. Follow these and you'll be well on your way to a secure wireless LAN:

- **Use WEP:** Even though it's full of holes, WEP will still prevent the casual passerby from trying to get to your stuff.
- **Change the default SSID:** Before an access point is put into production, this is the first thing that should be changed. In

addition, you should periodically change the SSID on all of your wireless equipment at regular intervals. Using the default SSID is bad for obvious reasons, and changing it every so often can help to thwart people who have figured out your old one, even though changing it can be inconvenient.

- **Use 802.1X for authentication:** While not perfect, 802.1X is better than WEP's authentication, although WEP will still be needed to handle certain authentications.
 - **Use secure tunnels:** Whenever possible, make use of software and services that provide end-to-end encryption such as VPNs, SSH, and SSL. Make sure to read the appropriate FAQs and properly harden your equipment to protect it from exploits.
 - **Carefully position antennas:** The less exposure to the outside, the less chance that your network will be stumbled across and used for nefarious, illegal, or otherwise inappropriate activities. To see the locations of some networks that have been "stumbled" upon, check out <http://www.netstumbler.com/>, which offers a comprehensive, updated database as well as software for locating these networks.
 - **Use filters:** While not effective as a stand-alone security measure, filtering so that only recognized MAC addresses are allowed access to your network can help to bolster the security of your wireless network.
 - **Segregate your wireless network:** This works best if you use a VPN connection from the wireless device. This consists of setting up your wireless access points outside a firewall and configuring that firewall for VPN access from the wireless devices.
- Although it is impossible to completely secure a wireless network, if you use the tips presented in this article, you can help to keep a majority of the attackers away from you. Remember: Most organizations with wireless networks have done very little to address security, and most attackers would rather go after those easy targets rather than go out of their way to get into your locked-down system. ☞

Think security when setting up an 802.11b wireless network

Apr 22, 2002

By Ron Nutter, MCSE, CNE, ASE

Many companies are already deploying wireless technologies, and others are only moments behind. But before your company implements an 802.11b wireless network, you should consider how you'd secure it. In this article, I'll show you some obvious and some not-so-obvious ways to keep your wireless network safe.

Permanent DHCP reservations

If you use DHCP with your wireless network, you may have reservations about someone hijacking an IP address and gaining access to your data. Permanent reservation in DHCP solves this problem by requiring the MAC address of the wireless card to make the connection between wireless card and access point. This DHCP reservation requires the MAC address and unique IP address of the wireless card. When you use only permanent reservations for DHCP IP assignment, the wireless card doesn't have to be configured any differently for your network than it would to be used on another network. The exception to this, of course, would be that you would have to configure the correct channel(s) to use, but this would depend on which card you're using.

How you configure your permanent DHCP reservations will depend on which operating system you're using on your DHCP server. For instance, in a Linux environment, the `/etc/dhcpd.conf` would be edited to map MAC addresses to IP addresses. On a Windows 2000 DHCP server, you would handle the configuration through the DHCP MMC.

For someone to hijack the IP address of your wireless network, he or she would have to override the MAC address of the card or have equipment to listen in on your network to see which MAC addresses or IP addresses are being used. If you need an even tighter lockdown on your wireless network, you can also use permanent reservations in conjunction with RADIUS accounting.

READ THE RADIUS RFC

For more information on RADIUS accounting, check out RFC2139 (<http://www.ietf.org/rfc/rfc2139.txt>).

Use a firewall between your wireless and wired networks

Though most networks have some type of firewall between the wired network and the Internet, many don't deploy firewalls between the wired network and the wireless network. Depending on the size of the wireless network, you may not need a firewall as sophisticated as what lies between your wired network and the Internet. The two features you'll want to put in place are port filtering and proxy server authentication.

With port filtering, you block some IP ports and allow others to pass. You should have two types of port filtering: static and stateful. Of the two, static filtering requires a more extensive setup, because you must define port usage going through the firewall in both directions. Stateful filtering is easier to set up, because you define port usage from only one direction, the side where the packet originated.

The trade-off in setting up stateful port filters is that there will be a little more processor overhead on the firewall. This occurs because the firewall has to build a table of the traffic going through the stateful filter. With this table in place, the firewall will know which traffic can pass through and which cannot.

To make things easier, when setting up port filtering, you should have some type of protocol analyzer to see the ports that are being used in the communication that you want to allow to pass. Since the wireless standard 802.11b is a little different than what is used on the wired portion of your network, you will need to use a different protocol analyzer. Two analyzers that work with wireless networks are the AiroPeek NX from

WildPackets.com and Sniffer Wireless from Network Associates. I've used the beta version of AiroPeek NX and have found it to be very simple to set up and use. You can also share the packet capture filters you set up in AiroPeek with its wired cousin EtherPeek. The sharing of packets between the two sniffers saves you from having to set up duplicate filters between products.

The second feature you should use with your firewall is a proxy server, the most common of which is HTTP proxy. With HTTP proxy, you can require users going through the proxy to authenticate before being allowed to pass through. Depending on what you are using for your HTTP proxy, the authentication screen will come up as either an HTML screen or Java applet. Using an HTTP proxy means you won't need to configure as many port filter exceptions for your Web traffic to pass through your firewall.

Depending on the type of firewall you use between the wired and wireless portions of your network, you may also want to consider a virtual private network (VPN) server. While it may seem like a bit of overkill to use a VPN on a local network, Wired Equivalent Privacy (WEP) as it ships with the wireless cards isn't totally private. Using a VPN server with it adds an additional layer of security.

Taking advantage of antennas

To get better wireless coverage in your building/campus and to make it a little more difficult for unwanted users to steal your wireless bandwidth, use directional antennas to focus coverage only where you need it. I've seen three types of access points: those with omnidirectional antennas, those that use the antenna in a wireless PCMCIA card, and those that don't come with any antenna.

The omnidirectional antenna distributes a signal over as uniform an area as possible, as shown in **Figure A**.

Directional antennas concentrate the signal to a specified location, as shown in **Figure B**.

When looking at antennas, consider which coverage pattern will work best for your company's needs. For example, if your antenna must be placed next to an external wall, your best bet will be the directional antenna.

External antennas vs. PCMCIA wireless cards with antennas


I've noticed something with access points that use the antenna in the PCMCIA wireless card that might cause you to consider using an external antenna (even if it is omnidirectional). During one installation, I found one access point that had a much weaker signal when using a different brand of wireless NIC.

Figure A




The omnidirectional antenna sends out a signal in all directions.

Figure B



The directional antenna sends out signals in only one direction.

However, when I used the same brand of NIC as the access point, the signal was much stronger. Since you may have wireless NICs being used by visitors/vendors, and thus they

may have different brands of NICs, you should consider using an external antenna to ensure consistent support. 

How to beef up wireless security

Apr 26, 2002

By Robert L. Bogue

Wireless connectivity is the panacea for many of today's network woes. It eliminates expensive cable runs and provides workers with more freedom: no more struggling with the short tether of a network cable. However, this freedom leaves many organizations worried about security. In this article, I'll review some security methods you can use to protect your wireless network. I will also discuss the weaknesses of these security solutions and provide some mechanisms to overcome these weaknesses.

Unauthorized access to your network

If you have no security established on your wireless network, it's easy for someone to set up a system and break in. If you have DHCP set up, someone can even get IP address information automatically. Without DHCP, the hacker can simply use a wireless packet sniffer to determine the IP addresses of the stations already on the network and pick one that's available.

One issue most organizations face is the false sense of security given by the corporate firewall. No matter how tight, big, or expensive the firewall is, it can't prevent wireless signals from getting into the hands of hackers. Firewalls are put in place to prevent intruders from gaining access to the internal systems. However, when someone drives up and logs on to the wireless network, there's typically no barrier between them and those sensitive internal systems.

Security options

To secure your wireless LAN, consider the following options:

- ▶ Service set identifier (SSID)
- ▶ Wired Equivalent Privacy (WEP) protocol
- ▶ VPN
- ▶ MAC restrictions

Service set identifier

The SSID is designed to allow two wireless LANs to operate in close proximity. The SSID is used on the client and the access point to bind their communications together.

If the SSIDs don't match between an access point and the network card, there is no communication between the two. Because of this, some administrators believe they can just change the SSID and no one will be able to access their wireless network. Since there's no SSID match, there's no risk of unauthorized users gaining access. Although changing the SSID is an important step in securing the wireless network, it alone does not guarantee the network's security.

To set the SSID on a Windows 2000 machine, open the Properties window of the network adapter. Click the Configure button and then select the Advanced tab. From the Advanced tab, select SSID from the Properties listing and enter the correct SSID in the Values field. Click OK and the SSID will be set.

IS YOUR SSID REALLY SET?

If you want to make sure your SSID is set, there are a variety of programs that allow you to search for and find wireless LANs. NetStumbler is one such program. These programs can interrogate the access points in the area to determine the SSID. Also, since the SSID is routinely transmitted on the wireless network, it can be observed with a wireless packet sniffer such as Sniffer Wireless.

Although there is no real weakness to overcome with the SSID, the point is simply to make sure that you keep these IDs private. Don't release them into the hands of anyone unless that person has a need for that knowledge (such as a member of the IT staff). If someone has the SSID of your access point, he or she is one step closer to breaking and entering.

Wired Equivalent Privacy protocol

The solution to prevent eavesdropping is encryption. Since security is so important for a wireless LAN, the adopted standard has been defined as an encryption mechanism supported by both access points and network cards. The WEP protocol supports two different key lengths: 40-bit and 128-bit. As with other encryption mechanisms, the longer the key the more secure the communication.

WEP will eliminate the ability for someone to walk up and listen to packets crossing your wireless network and will prevent such people from joining the network. Unfortunately, WEP isn't flawless; it can be cracked with the right tools. One such tool is AirSnort for Linux. AirSnort captures and simultaneously tries to crack the WEP key being used on a wireless network. According to statistical models, nearly five million packets must be transmitted across a network for tools like AirSnort to be able to crack WEP. The number of packets that will be on your wireless network in a given day varies substantially, but a busy wireless network could transmit more than five million packets a day. So in some cases, a hacker could use AirSnort to crack your WEP key and break in to your wireless network.

The biggest problem with this type of attack is that it can't be detected. The machine running AirSnort can be set up to not broadcast a single packet, so it's impossible to know that someone is listening to the network trying to determine the WEP key. Once the hacker has the WEP key, he or she can listen to all data transferred on the network and eventually join the network.

Another challenge of using WEP is that there's no common method of updating WEP keys all at once. Since WEP keys are required for every device, a change in the WEP key means that you must update every device. Because this is such a tedious, time-consuming process, it's rarely done, which means once WEP is cracked or if someone who knows the WEP key leaves the organization, that person will likely have access to the network forever. My advice? No matter how much time it takes, if you know your WEP key has been compromised, change it.

Virtual private networking

A better approach to securing traffic on your wireless network is to have wireless users connect to a VPN server behind the wireless network. The VPN server is also connected to the local network and can route traffic from authenticated users on the wireless LAN to the local network.

The setup for a VPN server is more difficult than utilizing WEP; however, IPSec and PPTP don't have the vulnerabilities that WEP does. IPSec and PPTP have both been used for quite some time, and no one has been able to break their encryption mechanisms, which makes the encryption provided by IPSec and PPTP secure, even in a wireless environment.

Additionally, a VPN server provides user-level authentication. This means you can control access to the network from each individual computer and on a user-by-user basis. For example, someone could steal a network card with a MAC address approved for use with the wireless network, but the person still couldn't access the network without a valid user name and password.

VPNs are more complex to set up than the standard wireless network, they add expenses to the network, and they require processing

time on the client workstations. Where WEP is implemented in the hardware of the network card, establishing a VPN requires your computer to perform the encryption manually.

In terms of complexity, a separate VPN server must be installed for use on the wireless network and all access points must remain on their own network. In most corporate environments, this would mean setting up virtual LANs (VLANs) on the existing switches. However, there are organizations that don't have switches that support VLANs deployed across the organization, so setting up the wireless network could require a new set of cable runs.

MAC restrictions

Another method of security for your wireless network is to restrict the access points so that they talk only to specific MAC addresses. While WEP and VPN technologies encrypt all the data packets traveling across the network, MAC restrictions are focused on allowing only certain trusted network cards to communicate to access points.

This additional layer of security is useful, but it has three primary limitations. First, the

access points must have the capability to turn on MAC restrictions. Second, you must have control of the cards that can access the network. Third, the list of wireless cards accessing the network must be small enough to fit within the limitations of the access point to store the addresses, or the access points must be capable of fetching the approved MAC addresses from a central database.

In most cases, MAC restrictions are used in conjunction with WEP or a VPN to provide a secondary layer of protection. MAC restrictions wouldn't be a good choice for an overall security solution.

A word of warning

In preparation for this article, I used the Mini Stumbler program on my Compaq IPaq Pocket PC. I found dozens of networks that were broadcasting their availability to the world. Just for clarity, I didn't drive out of my way to find these either. Most of them were sniffable from the local interstate. Approximately 80 percent of these networks didn't even use WEP to encrypt their data. If your company is serious about setting up a wireless network, consider using a VPN setup or purchasing a proprietary solution that can provide user-level security and an encryption mechanism that can't be easily cracked. ~

MAC ADDRESSES

Media Access Control (MAC) addresses are physical addresses assigned to each card. These addresses are unique to each card. On enterprise class access points, you can establish a list of trusted MAC addresses. Then, each access point will communicate only with cards that have a MAC address in their list.

Use WEP to improve security on your wireless network

Aug 27, 2002

By Laura Taylor

Wired Equivalent Privacy (WEP) is an optional IEEE 802.11 feature used to provide data confidentiality. In short, WEP is used to encrypt and decrypt data signals transmitted between Wireless LAN (WLAN) devices. WEP works by encrypting the wireless radio frequency between the access point and client device and is the minimum amount of security you should have enabled on your WLAN. If you don't implement WEP, hackers can obtain information about your wireless network through a sniffer trace and can then join it without your knowledge. Since your wireless Service Set-Identifier (SSID) is sent over the air in clear text, you need to use WEP to encrypt your data to protect it from hackers. WEP itself is not the strongest type of security you can implement on your wireless network, but it is one of the easiest ways to strengthen your wireless security network.

This article provides methods for using and configuring WEP on Cisco Aironet 350 Series Wireless LAN components. The Aironet 350 Series Wireless LAN product line is a set of wireless access devices that include access points and client adapters that can pass packets at speeds up to 11 Mbps. Before you can use and configure WEP, you'll need to install and configure the devices that use WEP.

NOTE

Though the Aironet 350 Series offers several wireless adapters, I'm going to use the PCMCIA adapter for the purpose of this lesson, since most of the time when you use a wireless network, you'll be using it on a laptop.

The Aironet 350 wireless LAN adapter has a list price of \$169 (PCMCIA version) for a single card, which includes all drivers for all platforms. This means that you can use the same card for Linux, Windows, Mac OS, or

even MS-DOS. The Windows platforms that it runs on include Windows 9x, Windows CE, Windows Me, Windows NT, Windows 2000, and Windows XP. All of the Windows platforms have slightly different installation and configuration procedures. I'll tell you how to set up the card for Windows 98.

Your task list

Here is a summary of the steps you'll need to follow to get your WEP-enabled adapter and access point up and running:

1. Install the wireless access point.
2. Configure the WEP security features of the access point.
3. Install the Cisco Aironet PC350 wireless LAN adapter device driver.
4. Configure and enable WEP for the adapter card on your laptop.

Install the access point

Your access point operates in the 2.4-GHz band, similar to how a cordless phone works. Like a cordless phone, your access point has an antenna on one side and a wired connection on the other. Your WEP-enabled client adapter talks to the antenna, which then sends the data through the wire to wherever it's headed. If it sounds simple, that's because it is.

First, you'll want to connect an RJ-45 Ethernet connector to the Ethernet port on the back of the access point. The Aironet PC350 should probably come bundled with an RJ-45 connector but it doesn't, so you'll have to purchase one separately if you don't already have one. Connect the other end of the Ethernet connector to your 10/100 Ethernet LAN.

A power adapter comes with the access point, and after you plug it into your electrical outlet, plug the connecting wire into the back of the access point. When you see the LEDs blink amber, red, and then green, you're juiced with power and ready to configure WEP.

Configure the WEP security features of the access point

Your access point comes with a default IP address of 10.0.0.1. You have to make your access point IP network and the network octet (or octets) in your laptop's IP address match. To do this, I recommend changing the IP address on your laptop to match a unique host address that is on the same network as your access point.

Figure A

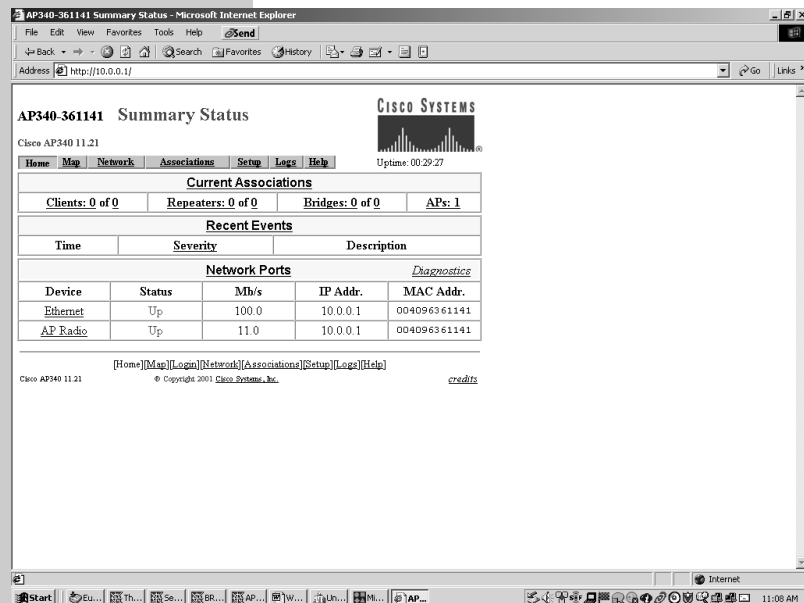
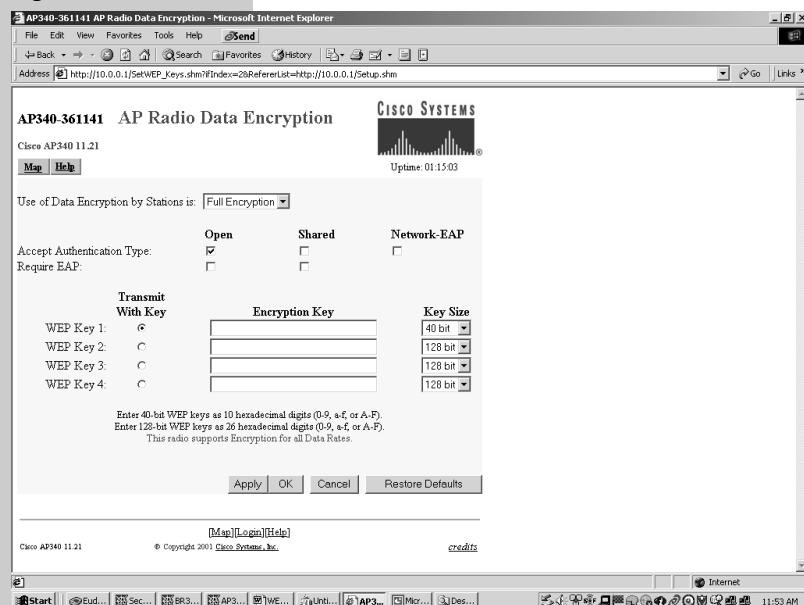


Figure B



The access point comes with a Web-enabled installation Wizard that guides you through the configuration process. Put the IP address of the access point in your Web browser like this: `http://<ipaddress/`

Don't forget to put the trailing slash on the end. Your browser will then map to a Web page on the access point that will display the Summary Status (see **Figure A**).

Once you are on the Summary Status page, you'll see various menu options on a row of buttons at the top. From the Summary Status page, click on the Setup button. Once you are on the Setup page, click on Hardware to see the AP Radio Hardware page.

The most important thing on the AP Radio Hardware page to fill in is a unique SSID. The SSID is a unique name you give to your access point. For the rest of the information on this page, you can just accept the defaults. On the bottom of the AP Radio Hardware page, click the link that says Radio Data Encryption.

The AP Radio Data Encryption page (see **Figure B**) is where you enter your WEP keys and select the key sizes. You can choose between 40- and 128-bit encryption for each key. How you set up WEP on your access point needs to match how you set up WEP on your adapter. You need to select one key for the transmit key (as a matter of best practice standards, select WEP Key 1 for the transmit key). Select 128-bit encryption because it is more secure, and it is very unlikely that you will notice any performance delays due to the higher encryption.

WEP on the access point for enterprise networks

If you're in an enterprise corporate environment, I suggest you enable broadcast key rotation for added security. Enabling broadcast key rotation eliminates the need to enter any keys in the boxes where it says Encryption Key, because the keys will be automatically generated. However, broadcast key rotation is available only if you use a RADIUS authentication server with Dynamic WEP keys. When you enable broadcast key rotation, the keys constantly rotate, making them much harder for hackers to sniff

with a protocol analyzer. To enable broadcast rotation, go back to the Setup page and then click on Advanced, which will take you to the AP Radio Advanced page (see **Figure C**).

If you go about halfway down the AP Radio Advanced page, you will see a dialog box where you enter a value for Broadcast WEP Key Rotation Interval in seconds. A zero value means the keys will not rotate. Keep in mind that the faster the keys rotate, the more potential there is for transmission latency while the key resets. I recommend starting with a small value, and if you encounter performance problems, increase it until the performance problems stop. A value of 300 would cause your keys to change every 5 minutes. Changing your key every 4 hours (14,400 seconds) is a good value to start with.

WEP on the access point for SOHO networks

If you are setting up WEP on a small office network, you should stick to static WEP keys. This means that you will not want to enable broadcast key rotation. However, you can change the keys manually, and you should change them at least once a week to decrease their accessibility to hackers.

There is a feature called Temporal Key Integrity Protocol (TKIP) that will add in extra security to compensate for not using broadcast key rotation. TKIP is a group of proprietary Cisco enhancements that include three methods of ensuring that your WEP keys cannot be cracked. One of the three TKIP features is Broadcast Key Rotation, which I mentioned earlier. The other two are Message Integrity Check (MIC) and Initialization Vector (IV) Hashing. SOHO users can turn on TKIP and MIC on the AP Radio Advanced page.

MIC prevents bit-flip attacks that occur when hackers intercept encrypted data and alter the bits slightly for the purpose of retransmitting them to destroy the integrity of the packet. IV Hashing modifies the headers of encrypted packets so that recurring patterns cannot be discovered or predicted by hackers. Along with regular WEP key rotation, these TKIP enhancements make WEP the most secure solution in wireless LANs today. Keep

Figure C

Figure D

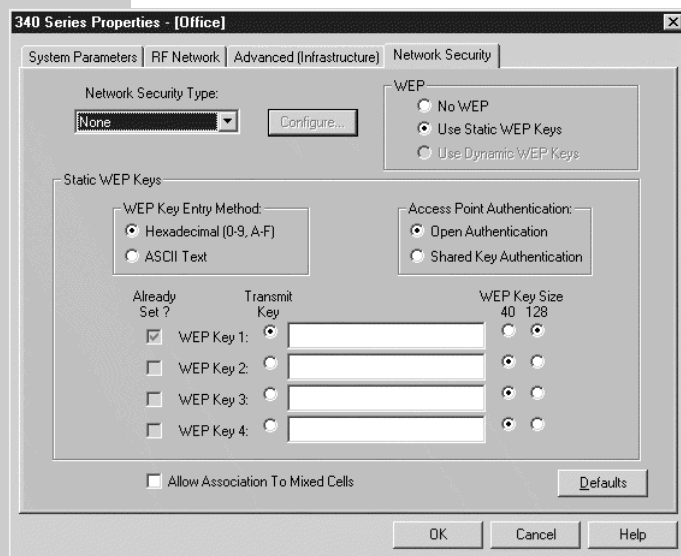
in mind that these are Cisco features that are above and beyond the 802.11b specification, so you will need a Cisco client card as well as a Cisco Access Point to enable these features.

Install the Cisco Aironet PC350 device driver adapter

After you insert the PCMCIA adapter into your laptop, Windows will automatically detect it, open the New Hardware Found window,

and collect information about it to build the driver information database. When you see the dialog box that says Windows Is Searching For New Drivers, click Next, and you will see a list of driver types. From that list, pick Network Adapters and click Next. Your wireless card is just another kind of network adapter. The next dialog box will ask you for the location of the driver, and you should select Have Disk. Insert the CD-ROM that came with your card, and Browse to the Win98 path on the CD-ROM drive. Now click OK. On the next screen, you should see the Cisco Wireless LAN Adapter already selected under Select Device, but in case it's not, select it and click OK. The Wizard will find the installation files and display the name of the client adapter.

Figure E



NOTE

Basically, each WEP key hex value is 4 bits, and each WEP key ASCII value is 8 bits. If you take the size of the WEP key and divide by 4 for hex, you get either 10 or 32 character values. If you divide by 8 for ASCII characters, then you get 5 or 16 characters for the key. What you use for a key doesn't matter, as long as the hex values range from 0-9 and a-f. With ASCII keys you can use any characters.

One final tip: There is an easy-to-use 128-bit Hex key generator at Leemon Baird's Web site (<http://www.leemon.com/crypto/MakePass.html>).

At this point, you may be prompted to enter the path to the Windows 98 operating system files. If the Windows 98 operating system files are already installed on your computer, put in this path name:

C:\Windows\Options\Cabs

If you are prompted for the Windows 98 operating system CD, insert the CD and browse to the proper CD-ROM drive letter and pathname, which in most cases will be D:\Win98. Now click OK. The required files will start copying to the proper location, and after this is complete, a dialog box will inform you that the Add New Hardware Wizard installation is complete. Click on Finish and reboot the computer to complete the process.

When the computer comes back up, select the Cisco Systems Wireless LAN Adapter and click Properties. Click on the Advanced tab and select Client Name. Type in your computer's name and then select SSID to type in your radio frequency (RF) network's SSID. Click OK to close the dialog box.

If you are using a static IP address, double-click My Computer | Control Panel | Network | TCP/IP Cisco Systems Wireless LAN Adapter. Click the Properties button, select Specify An IP Address, and enter the IP address, subnet mask, and default gateway address of the computer. Click OK. Then, in the Network window, click OK again, and you'll be prompted to restart your computer. When your system comes back up, your driver will be properly installed.

Configure and enable WEP for the adapter card

To configure WEP for the adapter card, you first need to get to the Series Properties screen by double-clicking the Aironet Client Utility (ACU) icon on the desktop. At the Series Properties screen, click on Edit, which brings you to the System Parameters screen (see **Figure D**).

In the client name field, enter your host-name and in the SSID1 field, enter the same SSID you entered when you configured the access point. Leave everything else on this screen as is, and click on the Network Security tab (see **Figure E**).

Tip

Do not accept the default SSID that comes with your adapter card. Having the correct SSID allows you to associate to the access point.

Network Security screen setup

Your first task is to decide if you want to allow communication with both WEP and non-WEP devices. Typically, in both enterprise and SOHO environments, you don't want to allow associations to Mixed Cells (the check box at the bottom of the screen), which means that you won't be letting wireless laptops communicate with non-WEP devices.


You'll also want to choose between the Open or Shared authentication options located in the Access Point Authentication pane. Open authentication means that users with the correct SSID will be able to associate to your access point; however, without the right WEP keys, their packets will be dropped. In Shared mode, both an encrypted and clear-text version of their data will be transmitted. Typically, Shared mode is preferable to Open since a user won't associate to the access point without the right WEP key anyway.

Finally, you need to decide whether you want to use static WEP keys or not. Static WEP keys don't change. Dynamic WEP keys are automatically generated and assigned to the adapter in a way similar to how DHCP automatically generates and assigns IP addresses.

If you are on an enterprise network and have thousands of wireless clients, assigning WEP keys can be quite a task. Enterprise users should enable broadcast key rotation on the access point, which means that the access point will use dynamic WEP keys. You want the adapter card and access card to work together, so select the radio button that says Dynamic WEP Keys in the WEP pane. Always select the first key as your transmit key (just as you did on the access point), and use the same level of encryption that you used on the access point.

If you are on a SOHO network you want to use static WEP keys. SOHO users should select the radio button that says Use Static WEP Keys and put in the same WEP Keys that were used for the access point.

Number generator tips

Static WEP keys can be generated in either hexadecimal or ASCII. For 40-bit keys, hex keys must be 10 characters long and ASCII keys must be five characters long. For 128-bit keys, hex keys must be 32 characters and ASCII keys must be 16 characters. Whatever value you put in for Key 1 on your access point has to match Key 1 on your adapter card. There is a very nice hexadecimal conversion chart at the Nickel Business Services Home Search Tools Web site (<http://www.nickeldesign.com/hexchart.htm>). 

Take steps to secure vulnerable WLANs

Oct 23, 2001

By Brian Hook

Unauthorized users may be lurking on your wireless local area network (WLAN), according to researchers at the University of California, Berkeley. The problem is caused by a number of key flaws in the Wired Equivalent Privacy (WEP) protocol, an algorithm that is supposed to protect wireless communication from eavesdropping and unauthorized access.

David Wagner, an assistant professor of computer science and a member of the WEP research team, said IT managers need to be concerned with a whole gamut of potential security problems posed by WLANs. Eavesdropping, tampering with transmitted messages, defeating access control measures, and denials of service are all potential threats.

Despite these security threats, wireless systems are becoming a hot commodity among businesses. Gartner released a study earlier this year forecasting that more than half of the Fortune 1000 companies will have deployed WLANs within two years.

With that in mind, here are methods you can use to secure WLANs in the face of these dangerous WEP vulnerabilities.

An easy hack

A wireless network uses radio waves to transmit data to everyone within range. So special precautions need to be taken to ensure that those signals cannot be intercepted. Wagner says his research shows that potential flaws in WEP seriously undermine the security of wireless LANs because hackers can easily break into wireless systems by using off-the-shelf equipment and positioning themselves within transmitting range of a WLAN. As a result, the WLAN is susceptible to a number of different types of attacks, including:

- ▶ Passive attacks to decrypt traffic based on statistical analysis.
- ▶ Active attacks to inject new traffic from unauthorized mobile stations based on known plain text.

- ▶ Active attacks to decrypt traffic based on tricking the access point.
- ▶ Dictionary-building attacks that, after an analysis of a day's worth of traffic, allow real-time automated decryption of all traffic.

WEP relies on a secret key that is shared between a mobile station and an access point. The secret key is used to encrypt packets before they are transmitted, and an integrity check is used to ensure that packets are not modified in transit. However, using the tactics mentioned above, it's easy to get around WEP. Wagner recommends that anyone using an 802.11b wireless network not rely solely on WEP for security. Instead, you should use other security measures to enhance WEP and WLAN security.

First step: Use WEP as the foundation

Despite the fact that he found major flaws in WEP, Wagner said it is very important that you use its encryption system as a foundation for good security.

"Surprisingly, a large proportion [of companies] deploy wireless networks without any encryption. So that is the first serious mistake that you can make," Wagner said.

"If you don't have WEP enabled—if you don't have [any] encryption enabled—[you are susceptible to] very serious attacks that require almost no sophistication. So the very first thing that you'd better do if you have a wireless network is...use encryption."

Second step: Isolate the WLAN and enhance encryption

After enabling WEP, you should also consider other security measures in order to compensate for its vulnerabilities. Wagner suggested a couple of steps to work around the potential problems of WEP.

"[First,] place your wireless network outside of the firewall. Treat it just like you

would the rest of the Internet,” Wagner said. “...recognize that it can’t be trusted and anything could happen on it, so you [should] fire-wall it off from all of your sensitive corporate secrets.”

Next, he said to use a virtual private network (VPN) for all traffic on the WLAN. The VPN will do its own end-to-end encryption on top of WEP. You can use such popular VPN protocols as PPTP and IPSec to accomplish this. Then, set up a VPN server/router that connects the WLAN segment to your LAN segment.

A cheaper, but less safe, alternative

Wagner admitted that the above solution might be too costly for some businesses, so he offered another suggestion that provides a limited defensive strategy.

First, it is important to understand that in WEP, there is a signal encryption key that’s configured identically for everyone who is supposed to have access to the wireless network. Usually this key is set up once when the password is handed out and often stays the same for months or years. That said, Wagner suggested that the wireless system employ extensions to WEP that perform dynamic key changes and modify the wireless encryption key once every 10 minutes.


“The problem is that once someone can break it, they’ve got everything,” Wagner said.

“So [by] changing the key once every 10 minutes, you can ensure that if they use this attack against you, they only get something that’s...10 minutes’ worth of data. And second of all, changing the key frequently makes it hard to mount [WEP] attacks.”

IT managers should be concerned

William Arbaugh, assistant professor of computer science at the University of Maryland, has also discovered flaws in WEP. He confirmed that WLANs are at great risk if they aren’t protected by additional security mechanisms.

“IT managers should be worried about unauthorized users accessing the corporate LAN via wireless access points,” Arbaugh said.

The research by Wagner and Arbaugh identified the risk posed by WLANs. Both researchers said it is wise to use WEP as a foundation but warned against relying on it as your sole method of security. Fortify it by placing your wireless network outside of your firewall and by using a VPN for all traffic and to connect the WLAN to your LAN. If that solution is beyond the scope of your budget, consider teaming up the WEP and dynamic key changes to protect your system. 

At last, real wireless LAN security

Aug 19, 2002

By George Ou

The freedom of wireless networking is enticing, but the accompanying risks are daunting. If you’re running a wireless LAN on the 802.11 standards, you may think your organization is secure. Think again. Joe User can drive to the local computer store, buy a wireless access point for less than \$100, and be free from Ethernet cables and any legitimate security within 15 minutes. And hunting

down one of these rogue access points is not an easy task.

The problem with WEP

During the inception of the 802.11 standards for wireless networking, the IEEE had to resolve a fundamental issue of wireless security: It’s vulnerable because it uses radio signals through open air space, as opposed to electrical

signals through closed wires. The Wired Equivalent Privacy (WEP) standard was created to address this liability. It was supposed to make wireless networks as private as wired networks by using 40-bit and 128-bit encryption. Maybe it's due to a lack of peer review or some other misstep, but whatever the reason, that "equivalent privacy" is not so private after all.

To be precise, WEP can be broken very quickly after gathering 100 MB to 1,000 MB of data with freeware sniffers commonly distributed on the Web. Anybody with a \$60 wireless PC card and a laptop can collect that data in three to 30 hours on a typical wireless network. From that point on, freeware utilities can easily break the WEP code.

Making things worse, range is not your friend—you're vulnerable to this type of intrusion from points way beyond your parking lot. Ten dollars' worth of stuff from Radio Shack and a Pringles potato chip can will boost an 802.11 card's 100-foot range to about a 10-mile line of sight. And we won't even discuss what an industrial-grade directional antenna can do to you.

Because the 802.11 standard has no facility to centrally manage or distribute keys, WEP is fatally crippled by the fact that its keys are the same for all users and all sessions, and the keys never change. Attempting to manually change the WEP keys is highly impractical.

Many IT pros think they've found an answer with the use of VPNs, but VPNs for wireless LANs are not very practical, convenient, or

totally secure. First of all, VPNs require users to take the extra step in making a VPN connection after securing a wireless LAN connection. In addition, any interruptions in service (which are common for wireless LANs) will terminate the VPN connection and force users to reconnect to the VPN server.

On the issue of security, only the traffic to the VPN server is encrypted, so the wireless LAN interface itself is left wide open, forcing the need to run a personal firewall on the WLAN interface. Many vendors have come up with solutions to address some of these security and convenience issues. But licensing is costly, and these products don't address the fundamental issue of wireless security. What is really needed is a WEP that *works*.

Introducing 802.1x and EAP

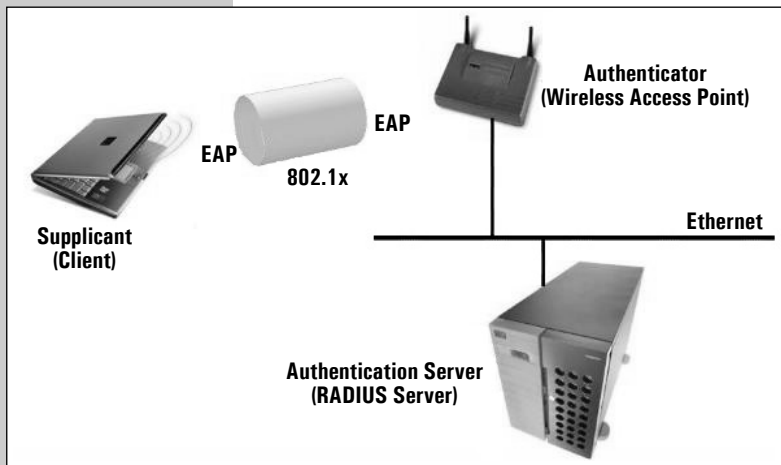
After the IEEE recognized the shortcomings of WEP and 802.11, it quickly came up with the 802.1x and EAP solution. A standard for Port-Based Access Control for both wired and wireless networking, 802.1x in itself does not make wireless networking secure. However, combine 802.1x with the Extensible Authentication Protocol (EAP) standard, and the gold standard in wireless network security is born; it's now possible to resolve WEP's biggest liability: static user and session keys.

User authentication is now mutually assured, and WEP keys can be centrally managed with policies and distributed securely. WEP keys can now be unique for individual users and individual sessions. In addition, keys can be set to automatically expire every 10 minutes to force constant rekeying, which makes it impossible to collect the 100 to 1,000 MB of data that hackers need to break WEP.

Figure A shows how this combination works.

The client makes a connection to the access point. At this point, the client is in an unauthorized state and not given an IP address or permitted access to the network in any way. The only thing the client can do is send 802.1x messages. The client sends user credentials to the access point with EAP, and the access point forwards the request to the Remote Authentication Dial-In User Service (RADIUS) server for approval. If the credentials are valid,

Figure A



the client will request credentials from the Authenticator via 802.1x and EAP. Once that process is complete, the RADIUS server issues a new temporary WEP key, and the access point allows the WEP session to proceed for that client. Every 10 minutes, the key expires and the EAP authentication process is run again to buy another 10 minutes of time.

Security is worth the investment

For any business network where wireless encryption needs to hold beyond one day, the

time for real wireless LAN security has arrived. It may cost a few times more than a consumer access point and require a more complex implementation, but your company's security should be worth a lot more than a \$100 SOHO wireless access point. Your \$100,000 firewall is useless if someone puts up a rogue access point, and standard WEP can do little to stop such attacks. ~

WPA wireless security offers multiple advantages over WEP

Aug 20, 2003

By Brien M. Posey, MCSE

For several years now, the primary security mechanism used between wireless access points and wireless clients has been WEP encryption. The problem is that although WEP encryption strength has increased a few times since Wi-Fi was introduced, the WEP protocol is still fundamentally weak because it uses a static encryption key. As a result, motivated attackers can easily crack WEP encryption by using freely available hacking tools.

Fortunately, some standard alternatives to WEP are emerging. The Institute of Electrical and Electronics Engineers (IEEE) has defined an expansion to the 802.11 protocol that will allow for increased security. Unfortunately, the standard is presently in draft form and isn't expected to be ratified until the end of 2003. In the meantime, though, most of the Wi-Fi manufacturers have agreed to use a temporary standard for enhanced security called Wi-Fi Protected Access (WPA). Although WPA is a temporary protocol and isn't recognized by IEEE, it is very similar to the revised IEEE standard expected by the end of the year.

Therefore, administrators that manage wireless LANs should become familiar with WPA.

802.1X authentication

If you have been using Wi-Fi for a while, you are probably familiar with the 802.1X authentication protocol. This protocol allows users to authenticate into a wireless network by means of a RADIUS Server. In standard Wi-Fi, 802.1X authentication is optional. However, 802.1X authentication is a requirement for WPA.

If your environment does not have a RADIUS server in place, you can still use WPA in spite of the 802.1X requirement. As an alternative to RADIUS, WPA supports the use of a preshared key.

WPA key management

One of the biggest drawbacks to traditional WEP security is that changing the encryption key is optional. Even if you do switch encryption keys from time to time, there is no option for globally rekeying all access points and all wireless NICs. Instead, rekeying is a tedious manual process and is completely impractical

for large organizations. After all, the instant you rekey an access point, none of the clients will be able to access it until they are also rekeyed.

But with WPA, the rekeying of global encryption keys is required. In the case of unicast traffic, the encryption key is changed after every frame using Temporary Key Integrity Protocol (TKIP). This protocol allows key changes to occur on a frame-by-frame basis and to be automatically synchronized between the access point and the wireless client. Global rekeying works by advertising the new keys to wireless clients.

The TKIP is really the heart and soul of WPA security. TKIP replaces WEP encryption. And although WEP is optional in standard Wi-Fi, TKIP is required in WPA. The TKIP encryption algorithm is stronger than the one used by WEP but works by using the same hardware-based calculation mechanisms WEP uses.

The TKIP protocol actually has several functions. First, it determines which encryption keys will be used and then verifies the client's security configuration. Second, it is responsible for changing the unicast encryption key for each frame. Finally, TKIP sets a unique starting key for each authenticated client that is using a preshared key.

Checksums and replay protection

When WEP was initially designed, IEEE took steps to ensure that an encrypted packet could not be tampered with. WEP-encrypted packets include a checksum value at the end of the packet. This value is a 32-bit code that is derived from the rest of the packet. The idea is that if something in the packet's payload changes, the checksum will not match the packet any longer and the packet can be assumed to be corrupt. This 32-bit code is called the Integrity Check Value (ICV).

Although ICV is a good idea, it just isn't secure. There are hacker tools that allow someone to modify a WEP-encrypted packet and to modify the ICV as well. By modifying the ICV to match the modified payload, the receiver will be unable to tell that the packet has been tampered with.

To counteract this type of hacking, WPA supports a security measure called Michael. Michael works similarly to ICV but calculates a Message Integrity Code (MIC) in addition to the ICV. The wireless devices calculate the MIC using the same mechanisms they would normally use to calculate the ICV.

The first major difference is that the MIC is only eight bits, as opposed to the ICV's 32 bits. WPA still uses an ICV in the same way that WEP does, but the MIC is inserted between the data portion of the frame and the ICV.

The MIC has two main purposes. First, it is encrypted along with the rest of the frame and makes it much more difficult to tamper with a frame's data. Second, the MIC contains a frame counter. This prevents someone from launching a wireless replay attack.

Implementing WPA

To take advantage of WPA, you must have adequate hardware and software. From a hardware standpoint, this means only that your wireless access points and your wireless NICs must recognize the WPA standard. Unfortunately, most hardware manufacturers won't support WPA through a firmware upgrade, so you may find yourself forced to buy new wireless hardware if you want to use WPA.


From a software standpoint, none of the Windows operating systems will support WPA by themselves. Windows machines with WPA-compliant hardware can use WPA, but only after you have installed the WPA client. The WPA client will work only for machines running Windows Server 2003 and Windows XP. You can download the necessary client from Microsoft (<http://microsoft.com/downloads/details.aspx?FamilyID=009D8425-CE2B-47A4-ABEC-274845DC9E91&displaylang=en>).

Mix and match

Obviously, switching wireless hardware and implementing WPA can be a big undertaking. Fortunately, it isn't something you have to do all at once. Wireless access points can support WPA and WEP at the same time. This allows for a gradual transition into WPA.

The only thing you need to know about mixing WEP and WPA is that doing so prevents

the global encryption key from being automatically rekeyed. Remember that WEP clients do not support automatic rekeying. To prevent key recognition problems, automatic rekeying is initiated by the access point only when no clients are running WEP. However, all of the other WPA security measures will work during the transition period.

As you look ahead to future WLAN deployments, keep in mind that you will probably want to change your security methods to encompass WPA and/or the similar set of security standards that is forthcoming from the IEEE. 

Six tips for implementing closed networking on a wireless network

Sep 30, 2002

By Scott Lowe, MCSE

Implementing a wireless networking system can result in serious security problems if the system is not properly secured. This is true of a wireless network deployed at home or one deployed in the office. In fact, some residential Internet service providers have clauses in their agreements that indicate that service is not to be shared with people outside of those covered by the agreement. If you deploy an insecure wireless network, it could result in a loss of service or in the use of your network as a launching pad for attacks against other networks. To help you close these security holes, here are six quick wireless networking tips.

Why do I want to close the loop?

The point of properly securing a wireless access point is to close off the network from outsiders who do not have authorization to use your services. A properly secured access point is said to be “closed” to outsiders. A wireless network is more difficult to secure than a typical wired network due to its nature. A wired network has a limited number of fixed physical points of access while a wireless network can be used at any point within the range of the antennas.

Plan antenna placement

The first step in implementing a closed wireless access point is to place the access point's

antenna in such a way that it limits how much the signal can reach areas outside the coverage area. Don't place the antenna near a window, as the glass does not block the signal. Ideally, your antenna will be placed in the center of the area you want covered with as little signal leaking outside the walls as possible. Of course, it's next to impossible to completely control this, so other measures need to be taken as well.

Use WEP

Wireless encryption protocol (WEP) is a standard method to encrypt traffic over a wireless network. While it has major weaknesses, it *is* useful in deterring casual hackers. Many wireless access point vendors ship their units with WEP disabled in order to make the product installation easier. This practice gives hackers immediate access to the traffic on a wireless network as soon as it goes into production since the data is directly readable with a wireless sniffer.

Change the SSID and disable its broadcast

The Service Set Identifier (SSID) is the identification string used by the wireless access point by which clients are able to initiate connections. This identifier is set by the manufacturer and

each one uses a default phrase, such as “101” for 3Com devices. Hackers that know these pass phrases can easily make unauthorized use of your wireless services. For each wireless access point you deploy, choose a unique and difficult-to-guess SSID, and, if possible, suppress the broadcast of this identifier out over the antenna so that your network is not broadcast for use. It will still be usable, but it won’t show up in a list of available networks.

Disable DHCP

At first, this may sound like a strange security tactic, but for wireless networks, it makes sense. With this step, hackers would be forced to decipher your IP address, subnet mask, and other required TCP/IP parameters. If a hacker is able to make use of your access point for whatever reason, he or she will still need to figure out your IP addressing as well.

Disable or modify SNMP settings

If your access point supports SNMP, either disable it or change both the public and private community strings. If you don’t take this step, hackers can use SNMP to gain important information about your network.

Use access lists

To further lock down your wireless network, implement an access list, if possible. Not all wireless access points support this feature, but if yours does, it will allow you to specify exactly what machines are allowed to connect to your access point. The access points that support this feature can sometimes use Trivial File Transfer Protocol (TFTP) to periodically download updated lists in order to prevent the administrative nightmare of having to sync these lists on every unit. ~

Don’t use MAC filtering as your only wireless network security solution

Nov 4, 2002

By William C. Schmied

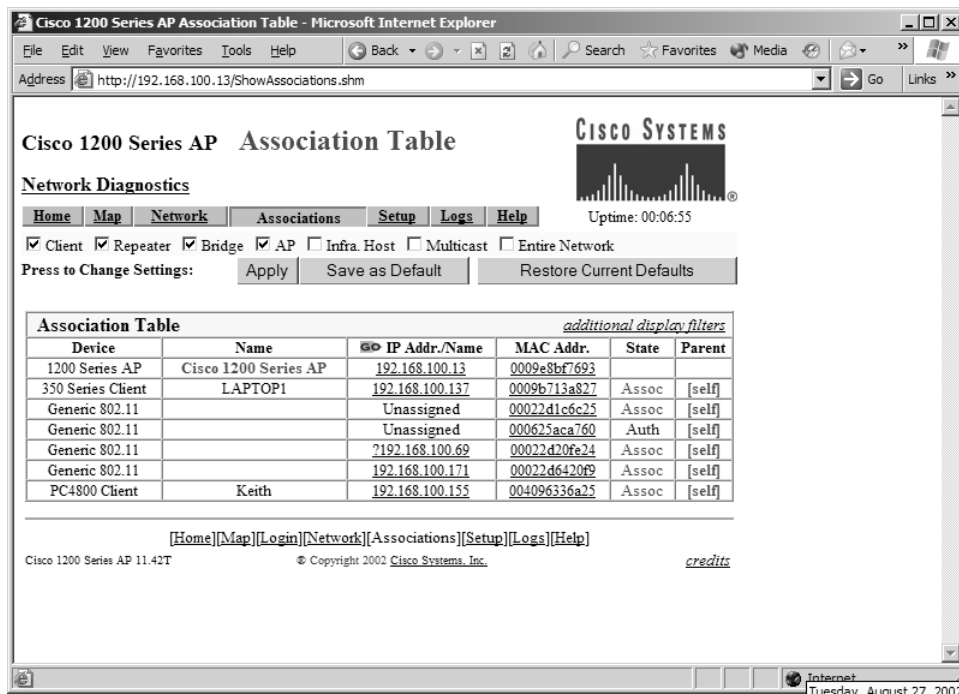
If you’re familiar with 802.11b wireless networking, you’ve no doubt heard the horror stories about how weak Wired Equivalent Privacy (WEP) is. In the rush to move away from WEP and its supposed weakness, many organizations have implemented Media Access Control (MAC) filtering as their sole wireless access point (WAP) security measure. What they may not know is that MAC filtering is extremely ineffective as a sole security measure. In reality, relying on MAC filtering to protect your wireless network is pretty much the same as leaving the front door open and asking an intruder to come on in and stay a while. In this article, I’ll show you how MAC filtering works and describe some of its pitfalls.

MAC filtering basics

Before I discuss why MAC filters aren’t the perfect security solution, let’s examine what MAC filters are and how they work. MAC filtering is the process of configuring an access point with a list of MAC addresses that will either be allowed or not allowed to gain access to the rest of the network via that WAP. The most common configuration has a list of allowed MAC addresses—the trusted and known MAC addresses that are supposed to be on the wireless LAN.

Exactly where you enter the allowed MAC addresses varies, depending on the WAP you use. Normally you’ll enter this information into the WAP’s configuration utility, usually

Figure A



Clients can be either authenticated or associated.

from a Web-based interface, although you can also do it from a console session or some other form of remote control. No matter how it's done, the end result is a list of MAC addresses that you use to allow or disallow access.

In **Figure A**, which was generated from a Cisco 1200 AP, you can see quite a few clients making connections to the WAP. Some are merely authenticated, while others are completely associated. In wireless-speak, “to authenticate to a WAP” simply means to announce your identity to the other station—in this case, the AP.

Authentication can take place using either open system or shared key (WEP) methods. To be associated with a WAP implies that the client is fully connected to the WAP and is now allowed to pass traffic through the AP. In short, the client now has complete access to the rest of the network, both wireless and wired. MAC filters act to keep unauthorized clients from becoming associated with the WAP.

An open door to intruders

The problem comes when an intruder wants to gain access to your network and has decided to

sniff your wireless network traffic. Sitting in your parking lot or some other easily accessible location, an intruder armed with the right hardware and software can easily sniff your wireless network and capture all packets sent to and from your access points. The captured data packets contain all the information the intruder needs to make a connection to your wireless LAN. This information includes the following:

- ▶ Authorized MAC addresses
- ▶ IP addresses
- ▶ IP subnets
- ▶ Wireless LAN SSIDs

The intruder can easily configure a wireless device with a captured IP address and subnet in the device's TCP/IP Properties window. Configuring captured SSIDs varies from one type of NIC to another, but it's done from within the configuration software provided with the NIC—again, a very easy configuration to make.

The tricky part comes in spoofing the MAC address itself. However, even an unskilled attacker can spoof a MAC address by making one quick registry edit. Using the

Registry Editor, all the attacker has to do is check the value of the NetworkAddress key, as shown in **Figure B**.

If the NetworkAddress string value doesn't already exist for the NIC or if it's blank, Windows reads the MAC address from the NIC's firmware. Entering a captured MAC address into the NetworkAddress string value for the rogue NIC tells Windows to use this MAC address for all communications emanating from the NIC. This registry setting works only if the NIC in the attacker's wireless device uses a PCI bus. This rules out most Flash Card-based NICs, but all PCMCIA cards, which appear on most laptops, use this bus.

After reconfiguring the rogue NIC with the stolen MAC address of an authorized client, the intruder will be able to seamlessly associate


with the WAP, which knows no different and is doing its job as it was configured to. If an attacker steals the MAC address during the day and doesn't use it until later—after the authorized user has left for the day—then the odds that the intruder will ever be caught are small.

Defense in depth

Just about all 802.11b access points support MAC filtering in addition to WEP. When used together, they form a pretty good security solution that will stop all but the most experienced and determined intruders. But MAC filtering alone won't cut it—even a relatively inexperienced attacker can get by it in 10 minutes or so.

So what do you do if you're responsible for a SOHO wireless network? You basically have two choices: 1) upgrade to wireless hardware that supports the Temporal Key Integrity Protocol (TKIP), which provides strengthening corrections for WEP, or 2) implement security by using both WEP and MAC filtering. For large, enterprise-level solutions, you should talk to your hardware vendor for a supported solution that increases your security. No matter what you do, don't go another day relying on only MAC filtering to keep intruders out of your network. ~

Figure B



Reconfigure the MAC address of the rogue NIC here.

Choosing a vendor solution for wireless LAN security with 802.1x and EAP

Aug 20, 2002

By George Ou

An emerging standard in wireless security finally is giving IT departments a way to fend off key-sniffing hackers and users who install their own unauthorized access points. In “At last, real wireless LAN security” (page 111), we discussed the new 802.1x/EAP combination that allows you to manage and distribute encryption keys on a user- and session-level basis.

Now we’ll tell you what it takes to actually build an 802.1x/EAP solution. Because 802.1x and EAP are open standards, implementation is left to individual vendors. As a result, four types of EAP implementations have emerged as “standards.” They all share the same underlying 802.1x and EAP architecture, but the ways they implement EAP are different.

LEAP

Cisco was one of the first vendors to market with its Lightweight EAP (LEAP) “standard” in December 2000. This is a very proprietary solution and initially worked only with Cisco Client 802.11 cards, RADIUS Servers, and Cisco Access Points. Recently, Cisco began working with other vendors to make its equipment and software LEAP-compliant. You now have some choice when choosing Client 802.11 PC cards, and at least four other RADIUS solutions support LEAP. Some laptop vendors even support this solution natively with their integrated 802.11 adapters.

Implementation of LEAP is relatively simple. Cisco’s ACS/AR RADIUS servers can easily be tied into your LDAP or Active Directory domain, and user authentication is transparent. The only downside to this approach is that your password policy better be strong, because LEAP is vulnerable to man-in-the-middle dictionary attacks. But if you have a strong password policy, LEAP is a fairly convenient and secure solution.

EAP-TLS

EAP-TLS (Transport Layer Security) is an open standard that’s supported by nearly every

vendor. As the most-common-denominator implementation of EAP, its strength is that it requires the use of public key infrastructure (PKI). PKI makes EAP-TLS extremely secure with the use of asymmetric public and private keys on the RADIUS and client sides.

The only downside is that implementing a PKI may seem a bit intimidating, although it really isn’t that difficult. Microsoft is firmly entrenched in this camp and has put native OS client support for EAP-TLS in Windows XP. Later this year, Microsoft will release support for Windows 2000, NT, 98, and Pocket PC. For the time being, you would have to use a third-party solution, such as that provided by Meetinghouse Data Communications (MDC), for non-XP operating systems.

Even Cisco is now recommending dual support for LEAP and EAP-TLS. EAP-TLS is a fallback solution with version 3 of Cisco ACS RADIUS because Cisco realizes that not everything is compatible with LEAP. The cost of implementing EAP-TLS is almost negligible if you use Microsoft RADIUS and PKI technology. This is because Microsoft’s Internet Authentication Service (IAS) RADIUS is bundled with the Windows 2000 Server operating system and is as stable as any other solution, in my experience.

Because Microsoft recommends that you implement IAS on your domain controllers, there’s no cost of an extra server and no additional licensing costs. The required PKI can be addressed by implementing the Certificate Authority (CA) service, also bundled with Windows 2000 Server. Licensing and server cost is kept to a minimum. Overall, this is one of the most secure and inexpensive solutions. The only initial burden is setting up a PKI in your organization; but keep in mind that PKI certificates can be used for many other purposes, such as L2TP VPN. All of this is just a one-time setup, and once EAP-TLS is fully implemented, it’s almost completely transparent to the user.

EAP-MD5

EAP-MD5 is the least secure version of EAP because it uses usernames and passwords for authentication and is vulnerable to dictionary attacks. In addition, EAP-MD5 does not support Dynamic WEP keys, which is a critical liability.

EAP-TTLS

EAP-TTLS (Tunneled Transport Layer Security) is Funk software's version of EAP that uses Funk's Odyssey or Steel Belted RADIUS server. It's also supported by third-party client software from vendors such as MDC. Funk's selling point is that PKI certificates are required only on the authentication server but not on the clients. In general, this is considered almost as secure as EAP-TLS while making deployment simpler.

PEAP PROPOSAL

Cisco, Microsoft, and RSA Security Inc. are currently proposing a new RFC for PEAP (Protected Extensible Authentication Protocol) to address the needs of organizations that want a more convenient password-based solution instead of the certificate-based solution used by EAP-TLS. Similar to EAP-TTLS, PEAP will require a certificate for the authentication server but not for the clients, and it will use an encrypted channel for password transmission to mitigate dictionary attacks.

Requirements for 802.1x and EAP

To use 802.1x and EAP, you must have the following components:

- ▶ Client wireless network adapter compatible with 802.1x
- ▶ Client access software capable of EAP
- ▶ Wireless access point (base station) compatible with 802.1x and EAP
- ▶ RADIUS compatible with EAP
- ▶ PKI

Most 802.11 wireless adapters support 802.1x natively with Windows XP. With older

operating systems, 802.1x driver support depends on the adapter's vendor. For Cisco LEAP-specific support, you'll most likely need to purchase a Cisco PC card. Very few 802.11 adapters support LEAP natively. Some of the Intersil Prism Wireless chipsets will support LEAP with the aid of third-party utilities. Some laptop vendors even have integrated 802.11 support for 802.1x and all four flavors of EAP, eliminating the need for bulky and expensive 802.11 cards. Most of the ORiNOCO adapters cost \$60 to \$100, while the Cisco adapters run between \$110 and \$140. Getting an integrated adapter from a laptop vendor with full EAP support will cost about \$50 to \$60.

For Client Access software, Windows XP provides OS native support for EAP-TLS. Microsoft will add support for older Windows operating systems, such as 2000, 98, NT, and Pocket PC, by the end of 2002. For LEAP support, Cisco's Client software was the only solution for some time. Third-party solutions such as that provided by MDC can offer EAP support for any of the four EAP types. Cisco's Client is bundled with its Wireless Adapters while some Integrated Wireless Solutions bundle the MDC solution.

For access points, only industrial-grade solutions will support 802.1x and EAP-TLS, such as those from Agere (a Lucent spin-off), Cisco, and Intel. However, LEAP currently works only on Cisco access points. These high-end access points cost between \$400 and \$1,000, depending on the features included. This is a bit more expensive than the SOHO solutions that cost between \$100 and \$200, but you get vastly superior features, including Dynamic WEP, better antennas, and sometimes even dual-band 802.11a and 802.11b capabilities.

For RADIUS capabilities, you can use FreeRADIUS on Linux (although support is shaky), Cisco's ACS/AR RADIUS, Funk Software's Odyssey or Steel Belted RADIUS, Interlink Networks, Open Systems Consultants, and Microsoft IAS (bundled with Windows 2000 Server). Pricing for the Linux and Microsoft Solutions are virtually free since you run IAS

on your existing domain controllers. The other solutions range between \$1,000 and \$4,000. It's important to note that all these RADIUS solutions support EAP-TLS. LEAP is supported by all but Microsoft. EAP-TTLS is supported only by Funk's solution.

PKI is required for the EAP-TLS and EAP-TTLS solutions. Microsoft Windows 2000 Server has the CA service bundled with the OS, so pricing is extremely attractive. Much of the PKI can be put onto your existing Windows 2000 servers. You can also purchase certificates from public CAs such as VeriSign, but that's not recommended for practicality and pricing issues.

As you can see, you have quite a few EAP choices, depending on your preferred platform. You can even bypass the EAP portion altogether if you go with Agere's proprietary AS2000 solution. But be warned that 802.1x and EAP will eventually be ratified into the 802.11i specifications. For most of you, the choice is between Cisco's LEAP (dominant

CISCO AND AGERE

While Cisco has a proprietary version of EAP, Agere uses its own proprietary encryption scheme, AS2000, that completely bypasses WEP and EAP while using 802.1x. However, both Cisco and Agere, like nearly all other vendors, support EAP-TLS.

market share), the standardized and super secure EAP-TLS solution with native server and client OS support, and Funk's EAP-TTLS. All have their own appeal.

The choice may be easier if you already are committed to many of the required components I listed. Just keep in mind that if you choose a proprietary solution, EAP-TLS should be implemented as a fallback solution for maximum compatibility. ~

Follow these steps to tighten security on Linksys wireless networks

Dec 10, 2002

By Lauri Elliott

By default, many wireless devices can leave networks and data open to access, paving the way for practices like war driving, in which someone armed with a wireless network card and a few easily obtainable hacker tools, can identify a wireless network and connect to it to access company data.

As network consultants, our mission is to provide the convenience of wireless networks in a relatively secure environment. To help you in this effort, here is a list of simple security fixes that will provide additional protection

when you're installing a Linksys wireless network access point for your clients.

Equipment used

The options I describe in this article will be based on use of:

- ▶ A Linksys wireless network access point; this device provides access for wireless clients to the wireless network.
- ▶ Linksys USB and PCMCIA network adapters for clients.
- ▶ A Windows XP operating system.

Stage one: Security configurations for the wireless network access point

In this first stage, you should make sure that the wireless network is running and clients are able to connect. You should note that some of the security configurations that I list here will make it more difficult to isolate network connectivity problems. But, ultimately, the enhanced security is worth the extra connectivity troubleshooting you might have to do down the road.

Figure A

The screenshot shows the Linksys Setup page. At the top, there are tabs for Setup, Password, Status, Log, Help, and Advanced. The Setup tab is active. Below the tabs, there is a message: "This screen contains all of the AP's basic setup functions. Most users will be able to use the AP's default settings without making any changes. If you require help during configuration, please see the user guide." The main content area is titled "SETUP". It includes fields for Firmware Version (1.01c), AP Name (enna-mi-001), and LAN IP Address (192.168.1.3). There are radio buttons for "Obtain an IP Address Automatically" and "Specify an IP Address". The Subnet Mask is 255.255.255.0 and the Gateway is 192.168.1.253. The Wireless section shows the SSID (rtfn458), Channel (6), and WEP settings (Mandatory, Disable, WEP Key Setting). The AP Mode section has radio buttons for Access Point, Access Point Client, Wireless Bridge, and Wireless Bridge - Point to MultiPoint. At the bottom, there are Backup and Restore buttons, and an Apply button.

Figure B

The screenshot shows the Linksys Log page. At the top, there are tabs for Setup, Password, Status, Log, Help, and Advanced. The Log tab is active. Below the tabs, there is a message: "Using this page to Enable or Disable logging." The main content area is titled "Log". It includes radio buttons for "Access Log" (Enable, Disable) and a "Send Log to" field (192.168.1.200). There are buttons for View Log, Apply, Cancel, and Help.

The configurations for stage one are:

1. Place wireless access point away from windows or exterior walls. The closer an access point is to a window or exterior wall the greater the signal will be outside the building.
2. Change the default settings for the access point. In particular, you should change the default IP address, the default service set identifier (SSID), and the default administrative password. To do so, access the Web-based administration utility on the access point, and then make appropriate changes to the Setup and Password pages. **Figure A** shows what you'll see, for example, on the Setup page.

Choose combinations that are complex for the SSID and password, which include letters, numbers, and special characters. The phrases should be at least nine characters long. Although this sounds like basic information, all too many businesses have neglected to perform this simple task and have found their networks compromised because of this oversight.

3. Enable logging. The log tells you which computers (by MAC address) have connected to the network. As with any log, you should do a quick scan on a daily basis to see if there is any unusual activity. To change the log, open the Log Web page within the administration utility. **Figure B** shows you what this screen looks like.

You can also have the log sent to another computer and view it using the Log Viewer utility provided by Linksys. I prefer this method because I can centralize my log files. Unfortunately, the Log Viewer is available only by sending an e-mail to Linksys Web site's support desk (<http://www.linksys.com/contact/contact.asp>).

Once you have completed these configurations, make sure all clients can connect successfully. You also should see what type of information is normally accessible by wireless network analyzers. A simple, free tool for this task is NetStumbler. **Figure C** highlights information accessible on a wireless network using NetStumbler.

Notice that NetStumbler identifies the access point, its maker, and the SSID. With this type of information, a person can connect to your wireless network. Therefore, it's now time to talk about how to reduce the likelihood that others will discover information about your network, connect to the network, and pull data from it.

Stage two: Security configurations

There are several methods for enhancing security on a wireless network. I'll examine a few of them.

Enable MAC filtering

With this method, you list the network adapters that are allowed to connect to the network by MAC address. The MAC address on a Linksys wireless network adapter is located on the bottom of the device. You can also get the MAC address by typing the command *ipconfig /all* (WINDOWS NT/2000/XP) at the command prompt while the wireless network adapter is installed on the computer.

The MAC address is listed as the Physical Address with this command. Once you have the MAC addresses, you can enable MAC filtering and list MAC addresses for clients you want to connect to the network. To access this page, you have to go to the Advanced tab in the Web-based administration utility for the access point (see **Figure D**).

Enable Wired Equivalent Protocol (WEP)

This method keeps outsiders from viewing data transmitted on your wireless network. Although WEP has come under fire because the protocol can be hacked, understand that your network is still more secure with WEP than without it. The key is to change the WEP encryption key regularly. I recommend doing it once a week, but many of you will feel this is too much work.

My advice is to balance the need for security with the administrative load. For those of you who are comfortable with scripting, you can create a script that will change the WEP passphrase (upon which the encryption keys are generated) and automatically update clients.

Figure C


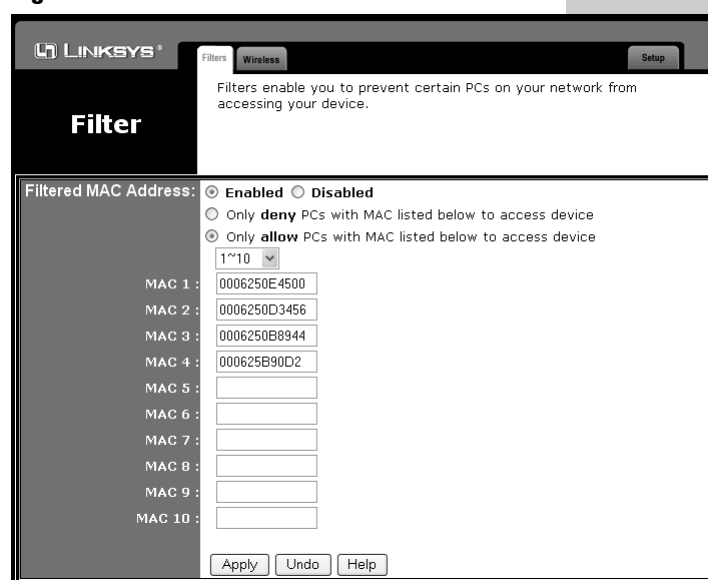


Figure D



More expensive wireless network equipment may have features built-in to do this. To set these features, you will use the Setup page to make WEP mandatory. Then use the WEP Setting page to generate the encryption keys in the Web-based administration utility for the access point.

Set encryption for 128-bit encryption. The higher the encryption, the more difficult it is to compromise it. Some wireless network devices provide 256-bit encryption as well, but both the access point and client network adapters need to support it.

Disable SSID broadcasting

Without the SSID being broadcast, your network is more difficult to locate. To set this

Figure E

LINKSYS®

Filters Wireless Setup

The advance Wireless Setting includes Beacon Interval, RTS Threshold, Fragmentation, DTIM interval, Rates, Authentication Type etc.

WIRELESS

Beacon Interval: 100 (msec, range: 1~1000, default: 100)

RTS Threshold: 2432 (range: 256~2432, default: 2432)

Fragmentation Threshold: 2346 (range: 256~2346, default: 2346, even number only)

DTIM Interval: 3 (range: 1~65535, default: 3)

Basic Rates: ☒ 1-2(Mbps) ☐ 1-2-5.5-11(Mbps)

Transmission Rates: ☐ 1-2(Mbps) ☒ 1-2-5.5-11(Mbps)

Preamble Type: ☐ Short Preamble ☒ Long Preamble

Authentication Type: ☐ Open System ☐ Shared Key ☒ Both

Antenna Selection: ☐ Left Antenna ☐ Right Antenna ☒ Diversity Antenna

SSID Broadcast: ☒ Enable ☐ Disable

Apply Cancel Help

option, go to the Wireless page under the Advanced tab in the Web-based administration utility for the access and choose Disable in the SSID Broadcast field (see **Figure E**).

Final check

After having done all of this, you can run Net-Stumbler again to see what type of information is accessible. You should find that none of your wireless network devices are located. Note that when WEP is enabled and SSID broadcasting remains enabled, the access point—including the MAC address—will still be visible; however, the name of the SSID will not appear. ~

XP client configuration for enhanced security on a Linksys wireless network

Jan 13, 2003

By Lauri Elliott

Chances are that some of your clients will be migrating from Windows 98 or 2000 to Windows XP this year. If your clients have a wireless network, you'll obviously want to take advantage of the security features offered in both the OS and the wireless network equipment. If, for example, you've configured a Linksys wireless network, the next step is to configure the Windows XP client—a topic I'll cover in this article.

NOTE

This article assumes that you have successfully installed the device driver for the Linksys network adapter and connected to the wireless network before applying the security enhancements.

Configure wireless network adapter in Windows XP

Because of Wired Equivalent Protocol (WEP), Windows XP's wireless zero configuration utility (WZC) will not be able to automatically connect the wireless network. Therefore, you will need to set some additional options in Windows XP. To make these changes, you'll need to:

- ▶ Double-click the network connection icon for the wireless network in your system tray on the desktop.
- ▶ Click the Advanced button at the bottom-left corner of the Wireless Network Connection dialog box (see **Figure A**).
- ▶ To add the wireless network as a preferred network, click the Add button in the


Preferred Network section. You'll then see the screen shown in **Figure B**.

- ▶ Type the service set identifier (SSID) for the wireless network in the Network Name (SSID) field.
- ▶ Check the Data Encryption (WEP Enabled) check box.
- ▶ Check the Network Authentication (Shared Mode) check box.
- ▶ Check The Key Is Provided For Me Automatically check box. (If you still have problems connecting to the wireless network, uncheck this option, then type in the first key generated by the WEP passphrase. You can get this information from the WEP Settings page in the Web-based administration utility for the access point.)

Problem locating the wireless network

Once you turn off the SSID broadcasting, clients might not be able to locate or connect to the wireless network. I discovered this prob-

Figure A



lem with Linksys PCMCIA network adapters (WPC11 version 3). Linksys says this happens because WZC does not support disabling SSID broadcasting. Therefore, this is a problem you might find with any Linksys network adapter that supports WZC.

Figure B

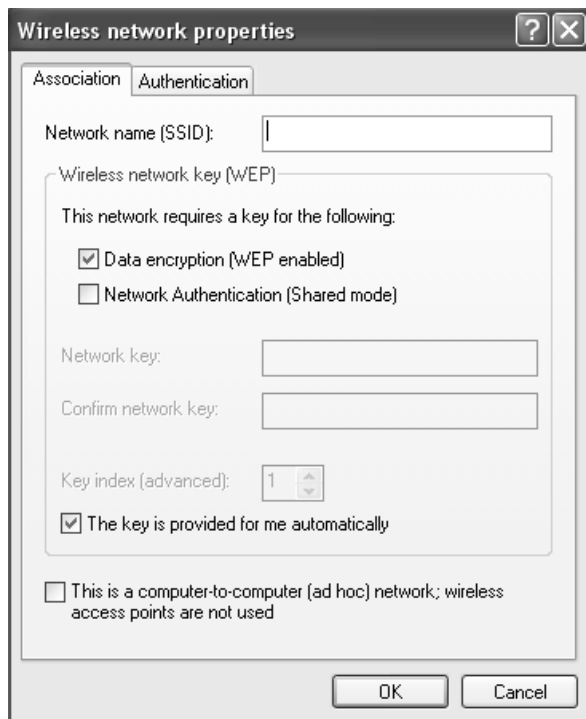
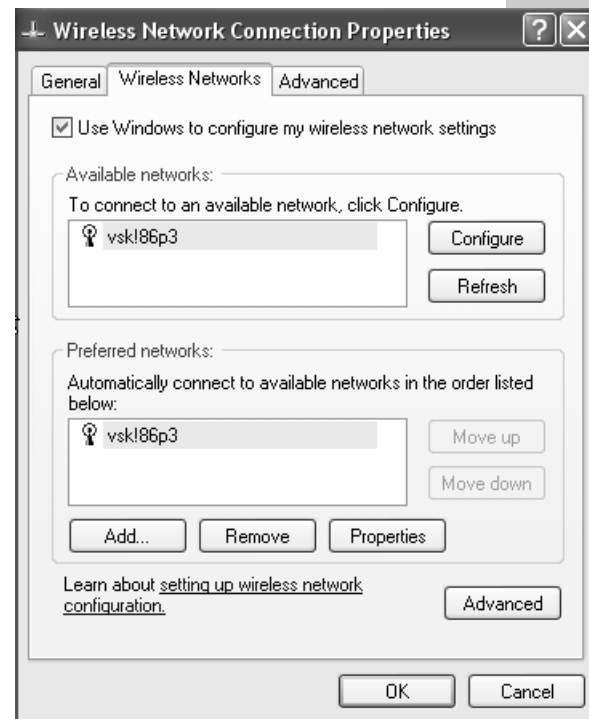


Figure C



Both Microsoft and Linksys indicate this is a problem, but they offer few workarounds. Linksys recommends that you use earlier versions of Linksys network adapters, e.g., WPC version 2.5, that do not support WZC.

Disable the WZC utility

The WZC service is not a requirement for a successful wireless network connection in XP. You can disable the service and get a slight improvement in system performance. To turn off the Windows XP WZC, do the following:

1. Right-click the My Network Places icon on your Windows desktop.
2. Choose the Properties option.
3. Right-click the network connection for the wireless network adapter.
4. Choose the Properties option.
5. Click the Wireless Networks tab.
6. Uncheck the Use Windows To Configure My Wireless Network Settings option (see **Figure C**).

You can turn off this feature entirely by disabling the WZC service in the Services Manager.

Device settings

When you disable WZC, you need to configure the wireless network connection options on the device profile. To access the device settings to be changed, as shown in **Figure D**, you need to:

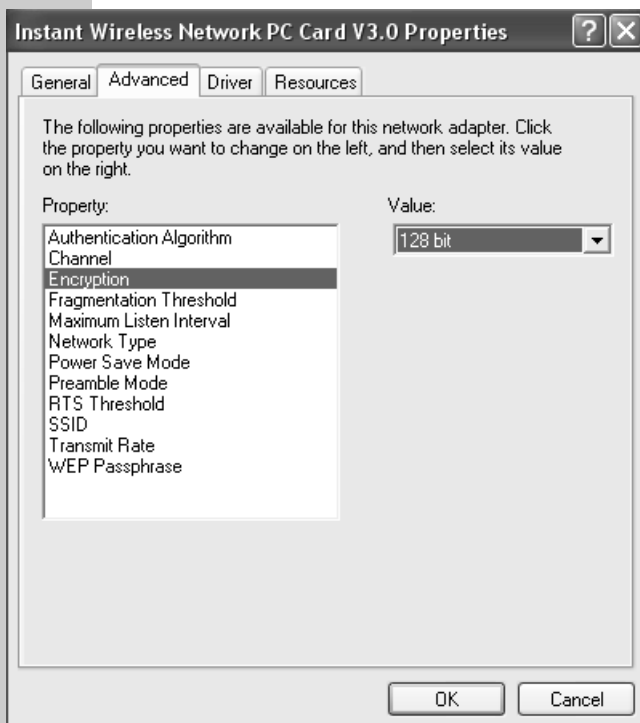
1. Right-click the My Computer icon on your Windows desktop.
2. Choose the Properties option.
3. Click the Hardware tab.
4. Click the Device Manager button in the Device Manager section.
5. Under the network adapters branch, right-click the Linksys wireless network adapter profile.
6. Choose the Properties option.
7. Click the Advanced tab.
8. Change the values for specific properties defined for the network adapter. The values to change are Encryption, SSID, and WEP Passphrase. Each should match the settings you defined on the access point.

Cutting down on the administrative headaches

In conjunction with these tips, there are a few techniques you can employ to reduce your share of administrative overhead:


1. Create a suborganizational unit just for Windows XP systems. (Windows XP has additional registry settings and policies that Windows 2000 does not.)
To make management cleaner, you can apply a group policy on an organizational unit that affects all computers for settings that would apply to all client computers in your environment. Then, create the suborganizational unit for Windows XP clients to manage only the XP-related settings. One of these settings can be for the WZC in the registry.
2. Customize an administrative template just for Windows XP systems. This administrative template will be attached to group policy on the suborganizational unit for just XP systems. To learn how to do this, check out the Microsoft TechNet article

Figure D



“Implementing Registry-Based Group Policy for Applications” (<http://www.microsoft.com/technet/treeview/default.asp?url=/TechNet/prodtechnol/windows2000serv/deploy/regappgp.asp>) and the Microsoft Knowledge Base article “HOW TO: Create Custom Administrative Templates in Windows 2000” (<http://support.microsoft.com/default.aspx?scid=kb;en-us;323639>).

3. Add the WZC service as an option in the customized administrative template. This will then be applied to all XP systems in the organizational unit.

4. Use AutoIt (<http://www.hiddensoft.com/>) to automate configuring device settings. This free application records keystrokes in Windows. You could record the keystrokes on the first system configured, then create a script to use with other systems. 

Notes

Notes

Products and Reviews

How to select the right wireless hardware for your home network	129
Go wireless with 802.11 options from Dell and Gateway	131
Supporting wireless users with 802.11 options from Compaq and IBM.....	132
Cut the cord with Agere Wireless USB Client systems.....	134
ORiNOCO's wireless network: Avoid its sticky setup problems	139
Installing ORiNOCO wireless gateway is a snap.....	142
ORiNOCO USB client setup makes a turn for the better.....	144
3Com AirConnect: Wireless for the great wide open	147
A review of 3Com's HomeConnect Home Wireless Gateway	154
Product Rating: 3Com Home Wireless Gateway	158
Connect wires and wireless with the Linksys Ethernet Bridge	160
Product Rating: Linksys EtherFast wireless AP and cable/DSL router with 4-port	164
Product Rating: NetGear MR314 cable/DSL wireless router	166
Product Rating: HP wireless gateway hn200w	167
Product Rating: Intel AnyPoint wireless gateway.....	169
Product Rating: SMC Barricade wireless broadband router.....	170
Product Rating: SMC EZ Connect 802.11a wireless access point	172
Quickly add wireless ports with SMC's EZ Connect wireless access point	174
Untether your network with SMC's wireless adapter	176
SMC's wireless broadband router offers performance tempered with caveats	178
Vivato's WLAN switches extend Wi-Fi range	182

How to select the right wireless hardware for your home network

Oct 16, 2002

By Greg Shultz

If you're thinking about setting up a wireless connection on your home network, you have some homework to do before you'll be ready to make the leap.

For starters, you have two choices, depending on whether you already have an existing broadband home network or whether you're starting from scratch. If you already have an existing home network and want to add a wireless connection to it, you'll just need a Wireless Access Point (WAP). If you're building a broadband home network from scratch and want to add wireless connections, you'll want to look at getting a wireless broadband router.

In this article, we'll examine each of these wireless options in detail. As we do, we'll help you decipher the technical terminology associated with wireless networking, and we'll discuss examples of some available products.

How wireless networking works

If you're new to wireless networking, you probably have a lot of questions, with the biggest one being how does it work. Well, wireless networking is simpler than you might think.

The easiest way to grasp the technology is to compare it to the little walkie-talkie sets that you played with as a kid. Each person has a portable handset unit with an antenna that allows each of them to wirelessly send and receive voice messages. You speak into a walkie-talkie, which converts your voice into radio waves and broadcasts them out over a small area. The antenna on the other walkie-talkie receives those broadcast radio waves and converts them back into voice waves that are then played on the walkie-talkie's speaker. The devices have a set range that allows users to communicate while they roam a relatively small area.

The principle is the same in wireless networking, except you're sending and receiving data instead of voice signals. Data travels to a wireless network device, where it is converted

into radio waves and broadcast to a relatively small area. Another wireless device receives those radio waves and converts them back into data.

On one end of the communication you have a base station, or WAP, that is physically attached to the network via a standard network cable. On the other end you have a wireless network card that will be connected to a desktop computer via a standard PCI slot, just like any other card. In the case of a laptop computer, the wireless network card could be in the form of a PCMCIA card with an attached antenna. Or the wireless network card could be built in to the laptop, with the antenna embedded into the lid on either side of the screen. Alternatively, you can also get wireless network adapters that attach to your computer's USB port.

Understanding wireless terminology

When you begin investigating wireless networking, you'll want to be familiar with wireless terminology, mainly in regard to wireless specifications or standards. These standards are designated by the number 802.11, along with a letter appended to it. The 802.11 designation is simply a number that was assigned to the wireless technology when the Institute of Electrical and Electronic Engineers (IEEE) began working on the project in the late '90s. Since the initial standard was finalized in June of 1997, several revisions have been made to the 802.11 standard, and that's where the letters come from. The revisions you'll run into when investigating wireless technology are *a* and *b*.

The 802.11b standard is the most common specification for consumer-oriented wireless products. It's also taken on a user-friendlier name, *WiFi*, which is short for Wireless Fidelity. In addition to the user-friendly name, the WiFi moniker specifies products that are

completely compatible with each other; a wireless product from one manufacturer that carries the WiFi logo will work with any other manufacturer's products that also carry the WiFi logo. Furthermore, WiFi products are priced within the reach of the average consumer (more on pricing in a moment).

Wireless networking products marked as 802.11b, or WiFi, work in the 2.4GHz band, have a maximum transmission speed of 11Mbps, and operate in a range of around 100 feet indoors. If the WAP has a direct line of sight to the outdoors, the operating range at 11Mbps jumps to 500 feet. Of course, you can still get a connection beyond 100/500 feet, but transmission speeds drop off.

The 802.11a standard is the most common specification for business-oriented wireless products. Wireless devices that use the 802.11a specification work in the 5GHz band and have a maximum transmission speed of 54Mbps. They also have an indoor operating range of around 300 feet and an outdoor range of a little over 1,100 feet. As you can imagine, the cost of 802.11a equipment is typically more than the average consumer wants to spend.

The stand-alone WAP

As we mentioned earlier, if you already have an existing home network and want to add wireless networking, you'll just need to add a WAP to your network. A WAP physically connects to your home network's hub or switch via a standard network cable. The WAP then

allows a computer equipped with a wireless network card to communicate with the network. You can get a basic WAP for anywhere from \$100 to \$200 from companies such as LinkSys, U.S. Robotics, 3COM, Belkin, D-Link, and Netgear.

The broadband router/WAP combination


If you're just starting out and are building a broadband home network from scratch, your best bet is to get a broadband router/WAP combination. As you can imagine, these devices provide all the same features as a broadband router and also function as a WAP. You can get a broadband router WAP combination for anywhere from \$100 to \$300 from leading vendors, including those we mentioned above.

The wireless network card

If you'll be connecting desktop PCs to your wireless network, you'll need to get wireless PCI network cards for each system. These range in price from \$10 to \$50. If you'll be connecting a laptop to your wireless network, you'll need to get a wireless PCMCIA network card. You can spend anywhere from \$25 to \$90 for this upgrade.

If you don't want to mess around with opening up your desktop PC or with PCMCIA cards for your laptop, you can get a USB wireless network adapter for around \$50.

If you're purchasing a new laptop to connect to your wireless network, you should ask the manufacturer about getting a built-in wireless adapter. Not only will this save you from having to insert and remove your wireless network device all the time, but you'll also find that the built-in antenna is more powerful than those used in the typical wireless networking PCMCIA card.

Depending on where you shop for your wireless networking devices, you may be able to track down a bundle deal that provides both the WAP and a set of wireless network cards. 

PROPRIETARY ENHANCEMENTS

When you're investigating home network-based wireless technology, you may encounter a new enhancement on the 802.11b standard being promoted by such manufacturers as D-Link and U.S. Robotics. This enhancement also operates 2.4GHz band but is being promoted as having a maximum transmission speed of 22Mbps. In addition to the increase in speed, there's an increase in the operating range that puts it in the realm of 802.11a. Now keep in mind that, at the time of this writing, this enhancement is not WiFi compatible. In other words, it's a proprietary technology that may not be compatible with wireless networking products from other manufacturers.

Go wireless with 802.11 options from Dell and Gateway

Aug 20, 2002

By Bill Detwiler, MCP

As the price of wireless networking equipment has dropped over the past year, the major computer manufacturers have begun offering customers a variety of wireless options. Whether your enterprise needs a hundred laptops with integrated wireless NICs or your small office needs a single access point and two wireless PC cards, two of the largest computer manufacturers, Dell and Gateway, have you covered. Let's take a look at several of the wireless options these two companies offer.

Dell's TrueMobile Wireless

Dell's TrueMobile customers have their choice of four wireless devices: an integrated TrueMobile 1150 Wireless Mini-PCI card, an external TrueMobile 1150 Wireless PCMCIA card, a TrueMobile 1170 Wireless Base Station, and a TrueMobile 1170 Wireless Access Point.

Wireless adapters

The integrated Mini-PCI card costs about \$100 and is available on several Inspiron and Latitude laptop models. The external 1150 PC Card, which retails for around \$70, works with all Inspiron and Latitude laptops and with all Dell desktops. To use the PC card with a desktop, however, you're required to purchase an additional PCI adapter card for around \$75.

Access points and base stations

The TrueMobile 1170 Wireless Base Station costs around \$175 and is designed primarily for the home or home office. It supports up to 16 wireless clients, allows for 128-bit encryption, and has a maximum open-environment range of 300 feet. The 1170 Wireless Base Station has a single 10/100 Fast Ethernet connection, can serve as an Internet router—when used with an existing cable or xDSL modem—and includes a NAT firewall.

For business environments, Dell offers the TrueMobile 1170 Access Point that retails for

about \$650. The 1170 Access Point is compatible with all 802.11b wireless cards; supports up to 32 wireless clients; has a built-in NAT firewall; works as a DHCP server; and has Ethernet, PCMCIA, and serial ports. The maximum range for the 1170 Access Point is 1,750 feet in an open office environment, 375 feet in a semi-open environment, and 165 feet in a closed office environment.

Gateway focuses on Intel equipment

Currently, Gateway customers who want to go wireless can choose equipment from both Proxim's ORiNOCO product line and Intel's PRO/Wireless line. However, because Proxim's equipment is being quickly phased out in favor of Intel's product line, I will cover only the Intel offerings here.

Access points and base stations

Gateway offers three different Intel access points: the Intel PRO/Wireless 5000 LAN Dual Access Point, the Intel PRO/Wireless 2011B LAN Access Point, and the Intel Wireless Base Station. Enterprise customers or those who are considering a move to 802.11a equipment should definitely consider the PRO/Wireless 5000 first. This dual-mode access point supports both the 802.11a and 802.11b protocols, offers 128-bit WEP encryption, and can handle up to 64 clients. The PRO/Wireless 5000 offers throughput speeds up to 54 Mbps for 802.11a clients, up to 11 Mbps for 802.11b clients, has a maximum range of 300 feet, and sells for around \$600.

For organizations that want a cheaper alternative to the PRO/Wireless 5000 and that don't need 802.11a support, Gateway offers the Intel PRO/Wireless 2011B LAN Access Point. This 802.11b-only device supports up to 60 clients, provides data throughput up to 11 Mbps, offers 128-bit WEP encryption, and has

a maximum range of 300 feet. Unfortunately, at \$450 the PRO/Wireless 2011B is only \$150 less than the PRO/Wireless 5000. If you can spare the extra \$150, I would definitely go with the PRO/Wireless 5000 with its dual-mode support and greater client capacity.

Intel's Wireless Base Station is Gateway's wireless option for the small or home office environment. Marketed primarily as an Internet connection-sharing tool, this 802.11b device features an integrated router, a NAT firewall, and a DHCP server. It can manage up to 16 wireless and 16 wired clients. The Wireless Base Station supports 128-bit encryption and costs around \$190.

Wireless adapters

When it comes to wireless adapters, Gateway gives customers a choice between Intel's PRO/Wireless 2011B LAN USB Device, Intel's PRO/Wireless 2011B LAN PC Card, and on some laptop models, an integrated

802.11b wireless adapter. All three of these wireless adapters support the 802.11b protocol, allow for data throughput of up to 11 Mbps, support 128-bit WEP encryption, and cost about \$90. The integrated wireless adapter is available on the 450 and 600 series laptops while the USB device and PC card can be purchased with any Gateway computer. However, there is a catch.

As I mentioned, if you want to use a wireless PC card with a desktop, you'll need a PCI adapter card. The Gateway sales representative that I spoke with said Gateway doesn't install PC cards and PCI adapters in their desktops. They would be happy, however, to sell me a PC card and a third-party PCI adapter that I could install myself. A wireless PCI adapter costs between \$35 and \$130, depending on the brand. ~

Supporting wireless users with 802.11 options from Compaq and IBM

*Sep 9, 2002
By Jeff Davis*

When it comes to supporting wireless networking, prices are going down and user expectations are going up. Whether you want to go wireless with four or five PCs in a small office or you need walk-around connectivity for a thousand corporate users, Compaq and IBM offer a number of options that can provide the proper access.

Compaq's wireless options defined by scope

Compaq defines its wireless solutions in terms of three distance-based platforms: the wireless personal area network (WPAN), the wireless local area network (WLAN), and the wireless wide area network (WWAN).

The WPAN connects devices that are relatively close to one another. The MultiPort lets you create a WPAN that connects Bluetooth-equipped devices with select Evo Notebooks, Presario Notebooks, and Evo desktops. In the iPAQ product line, the iPAQ Pocket PC H3870 model comes with integrated Bluetooth, and a Bluetooth Wireless Pack option is available for other iPAQ Pocket PC models.

If you need more horsepower in your home network, the Wireless Home Office Gateway model WL310 provides wireless connectivity to Internet service providers so multiple users can share the same Internet connection. It enables the sharing of common network resources such as files, printers, and scanners.

The WLAN solution provides access for small and medium-size businesses that need constant access across a couple of offices or a dozen floors in a skyscraper.

Some of the key components for Compaq's wireless solutions include:


- ▶ Compaq 802.11b-Bluetooth MultiPort Module for Compaq Evo Notebooks N400c and N600c. Priced at under \$200, it has a range of about 100 meters and supports TCP/IP, IPX/SPX, and UDP. The MultiPort Module relies on Bluetooth wireless technology to provide an interface between Bluetooth devices and Compaq notebooks.
- ▶ The 802.11b MultiPort Wireless LAN Module works with all Evo notebooks and keeps users connected to their network within a building, on a campus, or in a "hot spot" environment.
- ▶ Compaq's Wireless PC Card WL110 has a range of 525 feet in an open environment, 165 feet in a semi-open environment, and 80 feet in a closed environment.
- ▶ Compaq's WL215 provides high-speed wireless network for USB-enabled desktops or notebooks.
- ▶ The WL310 Wireless Home Office Gateway is compatible with Compaq Desktop, Compaq Notebook, and iPAQ Pocket PC. Its range is similar to the WL110's.
- ▶ The Compaq WL410 is a secure, full-featured access point that connects your Ethernet backbone and your wireless clients, supporting up to 50 users per cell.
- ▶ Wireless Enterprise Access Point WL510 is good for difficult-to-wire locations. It provides 10/100 Mbps Ethernet support over a wireless bridge.

IBM customizes wireless solutions by product lines

IBM's wireless accessories support ThinkPad notebooks or NetVista desktops using both Wi-Fi wireless and Bluetooth wireless technologies. As of this writing, the wireless product lines are organized in groups including IBM-specific, Cisco, Intel, and Novatel wireless. Some of the wireless devices available include:

- ▶ The IBM 11a Wireless LAN CardBus Adapter, which retails for \$179. It allows ThinkPad notebook users to stay connected in the office, home, or "campus" environment. It's compatible with the ThinkPad A Series, R Series, T Series, X Series, s Series, and i Series 1200/1300.
- ▶ The Bluetooth UltraPort Module For Bluetooth wireless communications, used for most ThinkPad and NetVista models, is priced at \$129. It allows a ThinkPad to communicate with other Bluetooth devices, so users can perform wireless tasks such as dialing out to the Internet using a cell phone or synchronizing with a PDA. Its range is around 30 feet (10 meters).
- ▶ The IBM High Rate Wireless Access Point 500, priced at \$449, provides wireless infrastructure for a home or small business that uses cable or DSL for Internet access. Using the popular IEEE 802.11b wireless technology, this gateway allows users to share a single cable or DSL account.
- ▶ The IBM High Rate Wireless LAN PC Card 128, priced at \$99, works in combination with the Wireless Access Point 500 to deliver the speed of IEEE 802.11b standards for wireless LANs and WLANs with the security of 128-bit Wired Equivalent Privacy (WEP) encryption.

If you're looking for a notebook computer designed for Cisco networks, check out the ThinkPad T30 notebooks. They feature integrated Cisco Aironet Wireless 802.11b. The Cisco Aironet 350 Access Point itself retails for \$569 and includes a 10/100 Ethernet uplink for integration with existing local area networks.

On the Intel side, IBM offers the Intel/Pro Wireless 5000 802.11a Access Point for around \$399. Its indoor range is 40 feet at 54 Mbps and 300 feet at 6 Mbps. Outdoors, the range is 100 feet at 54 Mbps and 1,000 feet at 6 Mbps. You can install up to eight access points per location to increase the available bandwidth to 432 Mbps and support more users. 

Cut the cord with Agere Wireless USB Client systems

Aug 2, 2002

By Jim Boyce

[Editor's note: Praxim, Inc., acquired the ORiNOCO product line from Agere Systems on Aug. 5, 2002.]

Networking might be a lot more common than it was 10 years ago, but it still isn't any easier to accomplish. As new technologies emerge, many of us in the trenches have had to brace ourselves—again—for the slippery slope of the learning curve. When you combine a new operating system such as Windows XP and a relatively new technology such as wireless networking, it can be an educational experience, to say the least. To help you make the transition, I'll explain the ins and outs of configuring a wireless network with Agere Systems' ORiNOCO Wireless USB Client.

Look Mom, no wires!

Agere Systems Inc. is a major player in the wireless networking market. Agere, which spun off as a separate company from Lucent Technologies (formerly AT&T) in June 2002, offers a range of wireless access products for everything from the backbone to the end user. The

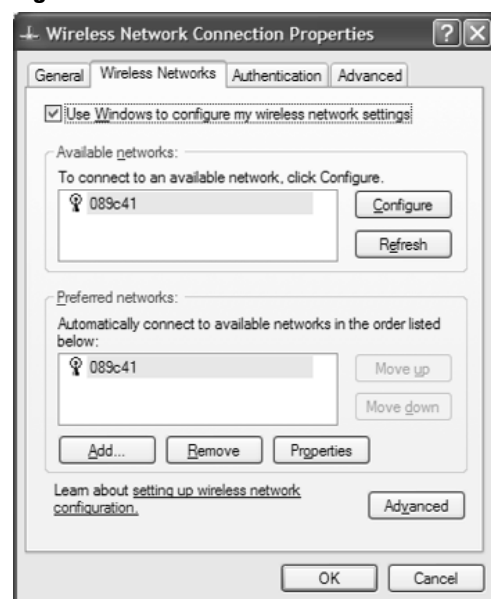
focus of this article is Agere's USB Client, but you can't really cover client configuration without looking in some detail at access points and related topics such as range and security. I'll start with a look at the USB Client and then work up the chain.

The ORiNOCO USB Client is an 11-Mbps desktop unit that connects to the computer's USB port with a standard USB cable. If you open up the USB Client unit, you'll find that the unit contains Agere's wireless PC Card with integral antenna and the hardware and firmware needed to adapt it to a USB connection. If you pop open one of Agere's access points, you'll find something similar: a wireless PC Card with the necessary support hardware and firmware. These are the same PC Cards you would use in a notebook's PC Card slot or in Agere's PCI or ISA adapters for desktop systems. By standardizing on a single unit, Agere not only simplified the product line but also cut its development and support costs, which should ultimately translate into better, less expensive products.

Choosing a USB Client over the PCI/ISA adapter means you don't need to worry about available slots or IRQs, and the unit doesn't need a separate power supply. Just one cable hooks it all up. The USB Client offers another advantage over the PCI/ISA implementation for workstations installed under a desk: In many cases you'll find that you need to add an optional antenna for PCI/ISA installations, particularly if the desk is metal. Using a USB Client lets you easily locate the unit on top of the desk where its range will be greater, which translates into better performance.

To connect to a wireless network, the client requires an access point. Like many of the wireless client products available today, Agere's USB Client can connect to any 802.11b (Wi-Fi) compliant access point or residential gateway. This includes Agere's Access Point, Access Server, and Residential

Figure A



Gateway products, as well as its older 2-Mbps and Turbo access point products. It also supports competing access points from companies such as Boingo, Nomadix, Linksys, D-Link, and others.

Actually connecting the unit is easy—after all, it has only one cable. Getting to the point of plugging in the cable takes a bit of setup, however, because you'll need to install the drivers for it before you connect the unit. Although Windows XP includes built-in drivers for the ORiNOCO wireless clients, you should still download and install the latest version rather than rely on the bundled drivers. You'll also need to download the software if you're installing the USB Client under other operating systems.

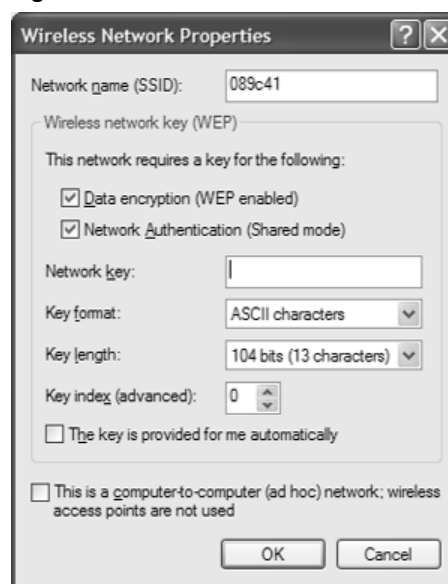
After you download the software, extract the file to a folder and run Setup.exe. Installation varies a bit depending on your operating system, but in all cases you won't have to do much besides specify the installation directory. When Setup is finished, you'll find a signal status icon on the system tray, which initially shows no device connected. Plug in the USB Client and wait a few seconds for the system to find the device. The system will tell you it has found a new networking device, install it, and then tell you it's ready to use. Now it's time to start configuring its settings.

Configuring network settings

The operating system determines the configuration options and methods you have for the USB Client. On all systems, Setup installs a Client Manager application you can use to view the card's status, choose a configuration profile, run diagnostics, and perform other configuration and testing tasks. Client Manager's Status area indicates the name of the connected network, signal strength, access point name, channel, and encryption status.

On Windows XP systems, Client Manager relies on the operating system's built-in wireless network configuration tools. On Windows 2000 and earlier systems, Client Manager provides its own wizard for configuring settings. Let's take a look at Windows XP first.

Figure B



Windows XP configuration

On a Windows XP system, you can open the Network Connections folder, right-click the Wireless connection, and choose Properties to open the Wireless Networks tab of the connection's properties sheet (see **Figure A**).

You can also navigate to the Wireless Networks tab by double-clicking the status icon on the tray to open the Client Manager status dialog box and then choosing Add/Edit Configuration Profile.

The Wireless Networks tab includes the following settings:

- ▶ **Use Windows To Configure My Wireless Network Settings:** This option allows Windows XP to automatically configure the wireless network. If you prefer to configure the settings yourself, clear this check box.
- ▶ **Available Networks:** This lists all of the available wireless networks detected by Windows XP. To configure the settings for a particular wireless network, select it here and click Configure.
- ▶ **Preferred Networks:** Where you have multiple wireless access points available, this list shows your connection preferences. Windows XP attempts connection to the network in the order in which the networks are listed. Use the Move Up and Move Down buttons to change the preference order.

► **Advanced:** Click Advanced to open the Advanced dialog box, which lets you set the following options:

- **Any Available Network (Access Point Preferred):** Use this option to allow Windows XP to connect to any available wireless network. It attempts connections to access point networks first. These are also called infrastructure networks. If Windows XP can't find an infrastructure network, it attempts a connection to an ad hoc network (computer-to-computer) if one is available.
- **Access Point (Infrastructure) Networks Only:** Select this option to prevent Windows XP from connecting to an ad hoc network if an infrastructure network isn't available.
- **Computer-To-Computer (Ad Hoc) Networks Only:** Select this option to prevent Windows XP from connecting to infrastructure networks, and to connect only to ad hoc networks. If no ad hoc networks are available, the connection fails.
- **Automatically Connect To Non-Preferred Networks:** Select this option if you want Windows XP to attempt a connection to a network even if it isn't listed in the Preferred Networks list.

To configure a specific wireless network, select it in the Available Networks list and click

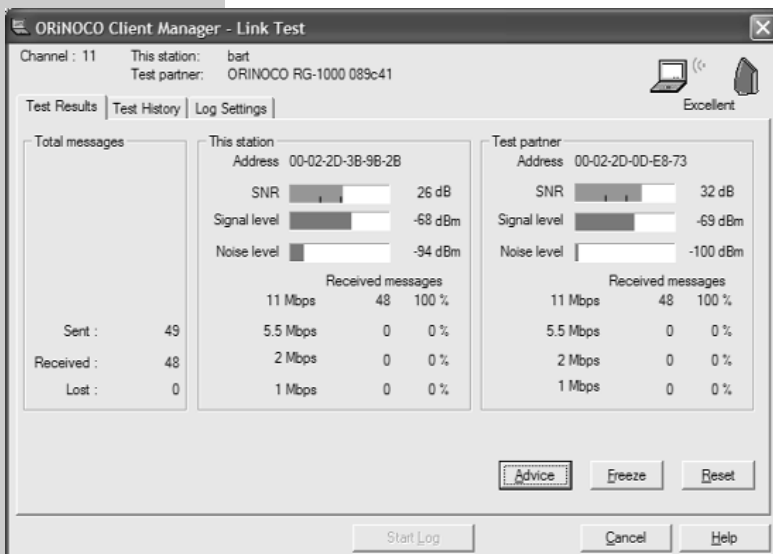
Configure to open the Wireless Network Properties dialog box (see **Figure B**).

The available settings are:

- **Network Name (SSID):** This field specifies the Service Set Identifier (SSID), which uniquely identifies the wireless network. In most cases, you won't change this value unless an automatically detected access point has changed names or you incorrectly entered the SSID for a manually added network.
- **Data Encryption (WEP Enabled):** Select this option to encrypt the data moving between the client and access point (or other wireless device). Your data is susceptible to interception without encryption, so you should enable this option.
- **Network Authentication (Shared Mode):** With this option selected (Shared Key Authentication mode), Windows XP uses the network key to authenticate the connection. With this option deselected (Open System mode), Windows XP does not authenticate the connection.
- **Network Key:** If you're not using a key provided by the device, enter the key in this field based on the key length specified by the Key Length field.
- **Key Format:** Use this option to select between ASCII and hexadecimal formats for the key.
- **Key Length:** Select the key length, either 40 bits or 104 bits.
- **Key Index (Advanced):** Use this spin control to select the location of the key.
- **The Key Is Provided For Me Automatically:** Choose this option to let Windows XP use the network key provided with the device.
- **This Is A Computer-To-Computer (Ad Hoc) Network:** Select this option to identify the connection as an ad hoc connection rather than an infrastructure (access point) connection.

If the access point or ad hoc connection you want to use doesn't show up in the Available Networks list, either the connection isn't available or it is configured not to

Figure C



broadcast its SSID. As I'll explain a bit later, hiding the SSID is a step you can take to secure your wireless network. In these cases, you can add the connection manually. To do so, click Add on the Wireless Networks tab to open a dialog box for the connection. Specify the SSID and other properties and click OK.

Other operating systems

Agere provides a wizard for configuring the USB client on operating systems other than Windows XP. To run the wizard, double-click the Client Manager icon on the system tray or open it from the Start menu. The Client Manager displays the current settings profile, named Default. You can modify this profile or create additional profiles for additional wireless networks. To modify a profile, choose Actions, Add/Edit Configuration File. Select the profile from the drop-down list and click Edit to start the configuration wizard.


In the wizard, you specify the profile name and choose a network type, whether it's access point, residential gateway, or peer-to-peer group (ad hoc). Next, specify the SSID for the network or click Scan to scan for available wireless networks. In the third wizard page, you enable or disable data encryption and specify the key and format. The fourth page lets you turn power management on or off for the client. The final page lets you configure the connection to renew its IP address when the profile is selected. You should enable this option if each of the wireless connections you use offers a different IP address range.

Running diagnostics

The Client Manager software includes some diagnostic tools you can use to test the card and monitor your wireless connections. You access these from the Advanced menu in the Client Manager. The Card Diagnostics perform several tests on the card, driver, and firmware. Keep in mind that testing the card temporarily disconnects it from the network.

The Link Test (see **Figure C**) provides a comprehensive look at signal and noise levels for the connection. The software identifies the stations by their MAC addresses; shows signal,

Figure D



noise, and SNR values; and offers several graphing options for analyzing the results. You can also configure logging for the connection and turn logging on and off.

The Site Monitor (see **Figure D**) gives you a means for analyzing properties for available wireless networks. These properties include MAC address, signal strength, noise level, channel, and others. Selecting the properties to monitor is as easy as selecting the items from a drop-down list. The Site Monitor is handy not only for identifying potential problems but also for monitoring multiple access points to find the one with the best performance.

Extending network range

One problem you will experience sooner or later is lack of signal strength. I've tested several devices from Agere and other manufacturers in different deployment scenarios and have had generally good results. In a few situations, I've had to make some changes to get enough signal strength to make the network usable. You can make some of these changes at the client side, but some need to be made at the access point. In other situations, you might have to make each change at both sides of the connection.

First, the USB client doesn't have to be tucked under the desk like a PCI/ISA

implementation. You can set the unit on the desktop or even on a shelf to get it away from power cords, monitors, speakers, and other components that generate EMF interference. If need be, add a longer USB cable to get the unit farther away from the computer and interference. The same holds true for the access points, so carefully consider placement when you install them.

As I've already mentioned, the ORiNOCO units use the same PC Card internally. The PC Card contains an integrated antenna, but the unit also has a jack for an optional external antenna, which Agere manufactures as well. In many situations, you can obtain better signal strength and therefore better performance by adding an antenna. The only situation in which you can do so without voiding your warranty, however, is when you're using the PC Card in an implementation where the antenna jack is readily accessible. For example, it's no problem to plug in the antenna if you're using the PC Card in a notebook computer or a PCI/ISA adapter. The reason you void the warranty in other situations is that you have to drill a hole in the case for the antenna wire.

To use an external antenna with the USB Client (remember the voided warranty), first unplug the unit from the computer. Then, grasp the unit by the base and gently pull off the cover. Locate the antenna jack in the end of the PC Card. Drill a hole in the top of the plastic cover of sufficient size to accommodate the antenna's plug. Pass the plug through the hole, plug it into the PC Card, and replace the cover. Plug the unit back into the computer to see if you've gained sufficient additional strength. Experiment with antenna placement as you monitor the signal strength, keeping the antenna wire away from power cords and other EMF-generating devices.

Check out your access points for optional antenna support. Many provide antenna jacks, but others—such as the ORiNOCO Residential Gateways (RG)—do not. The ORiNOCO RGs use the same PC Card as the USB Client, so you can use the same type of external

antenna with them as with the client. The RGs provide no external antenna connection on the case, so you'll have to open the unit and drill a hole in the cover to accommodate the antenna.

A final word on securing the network

No article on wireless would be complete without a few tips on security. There are several things you can and should do to secure your wireless networks. This is particularly important in a business setting where you have confidential information being transferred on the network, but it can also be important for another reason: keeping unauthorized users off your network. It's not impossible for an enterprising person in the business next door to gain access to your network and Internet connection if your network isn't properly secured.

One of the first steps to take is to change the default SSID for your access points to something that isn't easy to guess. Better yet, turn off SSID broadcast, which requires that the client know the SSID. This provides the benefit of better security, but unfortunately means users can't scan for the wireless connection. This, in turn, might mean more support calls to help new users find the connection.

Another step you can take on the access point side is to change the default community string for the access point's SNMP management. Many default to using the ubiquitous string *public*, so you should change it to an arbitrary string.

It's also important that you use data encryption for the wireless network as I discussed previously. Enable the Network Authentication option to provide additional security. Finally, you can decrease the chances for unauthorized users to gain access by not using DHCP for wireless connections. Specify static IPs for the clients, and choose an arbitrary subnet rather than the common 192.168.0.n or 10.0.0.n networks. Keep in mind that this strategy becomes impractical as the number of clients grows, but it can be effective for small networks. ~

ORiNOCO's wireless network: Avoid its sticky setup problems

Jul 16, 2001

By Mike Walton

Setting up a wireless network shouldn't tax your problem-solving skills, but that could be exactly what happens if you opt for the ORiNOCO solution.

Figuring out the setup problems is well worth it in this case, as the ORiNOCO solution, once it's installed, functions as well as its advertisements claim.

In this review, we will fill you in on ORiNOCO's quirky setup software as well as some of the lacking documentation, which understates some of the most important elements of the installation. After reading these pointers on the ORiNOCO business solution, you should be able to perform a hassle-free installation for your enterprise.

The equipment we'll discuss here is:

- ▶ ORiNOCO Access Point 500
- ▶ ORiNOCO Silver PC Card
- ▶ ORiNOCO PCI Adapter
- ▶ ORiNOCO USB Client

Our wireless network

We installed all of this equipment using Windows 2000 Professional on Dell and Hewlett-Packard computers that were connected to a test network. Our building is a particularly difficult environment for wireless technologies because its construction is such that cell phones typically lose service connections within the walls.

Even with the constraint of an unfriendly building, once installed, the ORiNOCO equipment was able to see the access point from various distances. The farthest point before the signal was reduced significantly was approximately 140 feet. The signal depreciated even more until it finally died at about 280 feet.

When the quality was still in the acceptable range, network traffic stayed close to 11 Mbps, but when the laptop got more than 200 feet

away from the access point, it dropped to around 5 Mbps.

While the equipment performed well once it was installed, getting to that point took some work. Much of our aggravation stemmed from a lack of experience with this hardware and software. To avoid this during your installation, make sure you have the most up-to-date drivers for your systems. We ended up with three or four different install CD-ROMs, and while they seem to be updated on a regular basis, if the date is older than six months, it may be worth a trip to the ORiNOCO Web site to download the updates.

One other general oddity that needs mentioning is that we were never able to get the Access Point 500 to access our test network through a 3Com hub. It would work fine when directly plugged in to the test network, but not being able to use the hub complicated our setup and implementation. After completing this review, we found that our problem was because the port on the Access Point 500 is 10 Mbps and the particular hub we were using was only 100 Mbps.

Figure A



The ORiNOCO Access Point 500 connects the wired and wireless networks.

Of course, in all the client computers, the network properties had to be configured after the client device (card or USB client) was installed.

Accessing the Access Point

The ORiNOCO Access Point 500 is a small beige box that contains an ORiNOCO Gold PC Card on a circuit board with an AMD computer chip.

On the bottom of the box is an RJ-45 port, a power cable port, and two inset buttons for

resetting or reloading the configuration on the access point (see **Figure A**).

The instructions indicate that when the access point powers up, it should have a set of default settings that will allow you to access the configuration file in order to customize settings, such as its IP address, from the included AP Manager tool.

Configuring the Access Point 500 from a wired network requires that the AP Manager tool be installed on the administration machine. Our situation was complicated by the fact that the access point would not go through our hub.

The obvious solution was to stick a PCMCIA card—the ORiNOCO Silver PC Card—into a laptop along with the AP Manager program to access the Access Point 500 (see **Figure B**).

To install the PC Card, the documentation recommends installing the Client Manager software first and then inserting the PC Card into the PCMCIA slot on the laptop.

Before inserting the card, we installed the Windows 2000 driver, which is a separate button under the Install Software portion of the Install CD. When we inserted the card, the Plug-and-Play feature in Windows asked us to find the driver, which we did by browsing our hard drive until we found the driver in the ORiNOCO folder inside the Program folder.

Hint: If the Wireless Network Control Panel applet doesn't show up in the Control Panel, you won't be able to configure your PC Card to find the access point.

We had to uninstall the driver and reinstall it before we were able to access the configuration file for the PCMCIA card.

Another hint: The Network Name on the PC Card and the Access Point 500 must be the same, or they will never be able to talk to each other. By default, they are supposed to be the same at original startup, but ours differed.

Our access point was essentially isolated because we couldn't get to it via wireless connection or through the wired network because the hub was preventing a network connection.

Figure B



The ORiNOCO Silver PC Card provides the connection to the access point.

Figure C



Compared to the rest of the installation, the ORiNOCO USB Client was a breeze to set up.

Our solution was to read the AP Manager documentation, which described setting up the access point through a direct cable connection through a hub. As the hub was not working for us, we ended up using a crossover cable directly between an onboard NIC on the laptop and the access point.

In the AP Manager, it is helpful to know that if there is any doubt about the settings of the access point, you can download the access point's configuration file under the Access Point menu. It allows you to save the file as another name and then view the contents of the file by going under the File menu to Edit Local Config File.

From the copy of the access point configuration file, you can then make sure you use the same network name for any other client access configuration files.

Setting up other clients

Along with the laptop and PCMCIA client combination, we set up both a USB client and a PCI adapter client on our wireless network.

The ORiNOCO USB Client was the easiest device on the network to install (see **Figure C**).

After installing the Windows 2000 driver and the Client Manager program from the CD-ROM, the Add New Hardware sequence began after we plugged in the USB Client. The instructions indicate that the driver is on the CD-ROM, but we found that we had to browse to the machine's hard drive to find the driver that we installed earlier.

The PCI Adapter was as difficult to install as the USB Client was easy (see **Figure D**).

The first thing we did was install the Adapter card in a PCI slot on the computer. Then we started the machine, went to Add/Remove Hardware in the Control Panel, and added the PCI-1410 CardBus Controller.

Next we installed the Windows 2000 driver and then plugged a PCMCIA card into the adapter. After doing this, the Add New Hard-

Figure D



More problems cropped up after the installation of the PCI Adapter.

NOTE

If you are installing the PCI card in a Windows 2000 machine, ORiNOCO warns that you must first make sure that you have updated to Service Pack 1 or higher.

ware sequence should begin, the appropriate driver should be selected, and then the Client Manager properties are supposed to be set.

Again, we had to uninstall and install the driver to get the Control Panel applet to show up so that it could be configured. Then we went to set the network properties on our new wireless NIC.

Surprise! When we hit OK on the properties, our computer rebooted immediately. It did this several times before we discovered on the ORiNOCO Web site a registry hack that needed to be downloaded and installed.

We then uninstalled and reinstalled an updated driver for Windows 2000 that we downloaded at the same time.

Only then were we able to set the Network Name for our PCMCIA card to find the Access Point 500. ~

Installing ORiNOCO wireless gateway is a snap

Sep 3, 2001

By Mike Walton

Wireless networking can be a great option when wiring CAT-5 cable would be difficult or impossible, coverage is limited, or the office is in a historic building that can't be modified.

Of course, even if you're not faced with any of these problems, wireless networking can provide employees with the option of moving around the building with their laptops while being connected to the Internet or the local network—an obvious advantage for some industries.

With all this in mind, the ORiNOCO Residential Gateway (RG-1000) is worth looking into as a wireless solution. The icing on the cake, especially for small or home office networks, is that the RG-1000 is a snap to install—as long as you are using installation software the company released in Summer 2001 or later.

Here's what we found when we set this product up on our test network.

Out of the box and online

Our tests of the ORiNOCO Residential Gateway were done on our test network using the gateway standing alone in one office; the setup of the gateway via the test network from another office; and the client, an ORiNOCO USB Client Gold, plugged in to a test machine in an office about 90 feet down the hall of a building notoriously unkind to wireless radio waves.

We have already reviewed several other pieces of ORiNOCO wireless equipment in "ORiNOCO's wireless network: Avoid its sticky setup problems" (page 139). As the title suggests, our biggest complaint about the ORiNOCO wireless solution involved the needless complexity of the software used to get the network up and running.

Since that review, however, ORiNOCO released new setup software that has reduced installation time to mere minutes—as long as it takes to open the box and plug everything in.

As with the other ORiNOCO equipment reviewed before, our test machine had a good 11 Mbps connection between the client and access point, which we tested out by playing a little Counter Strike on the Internet.

What you get in the box

We found the ORiNOCO RG-1000 selling for about \$280, and if you get the Start-Up Kit, which is what we tested, you get an ORiNOCO PC Card included in the box for another \$90.

The RG-1000 features a number of other handy capabilities, particularly for small or home office use. These include:

- ▶ An onboard 56-K V.90 modem with RJ-11 jack for telephone connections to an ISP
- ▶ A 10Base-T Ethernet connection through an RJ-45 connector

Figure A



On the left is the ORiNOCO Residential Gateway Start-Up kit, which comes with a PC Card and access point. On the right is the ORiNOCO USB Client Gold used in this review.

- ▶ Network address translation (NAT) for allowing up to 10 computers to use the same Internet connection
- ▶ A selection of four frequency channels
- ▶ 64-bit WEP encryption
- ▶ The ability to use a static or DHCP address for the network connection
- ▶ Reset and reload buttons for troubleshooting purposes

Getting it going

Getting the wireless network up and running using the RG-1000 could hardly have been simpler.

Once out of the box (see **Figure A**), you open the back cover of the access point and plug in the power supply and your connection to either a telephone line via the built-in modem or, as in our case, an RJ-45 CAT-5 cable connected to our test network.

Make sure that your CAT-5 cable plugs in to a switch or hub that will work at 10 Mbps, if you use that option. This is rarely a problem in the home office environment that is using a cable modem or DSL.

As far as software is concerned, you will need to install the setup utility on the machine you will use for that purpose. If you are accessing the access point via wireless connection, you will need to install the client software to access the RG-1000.

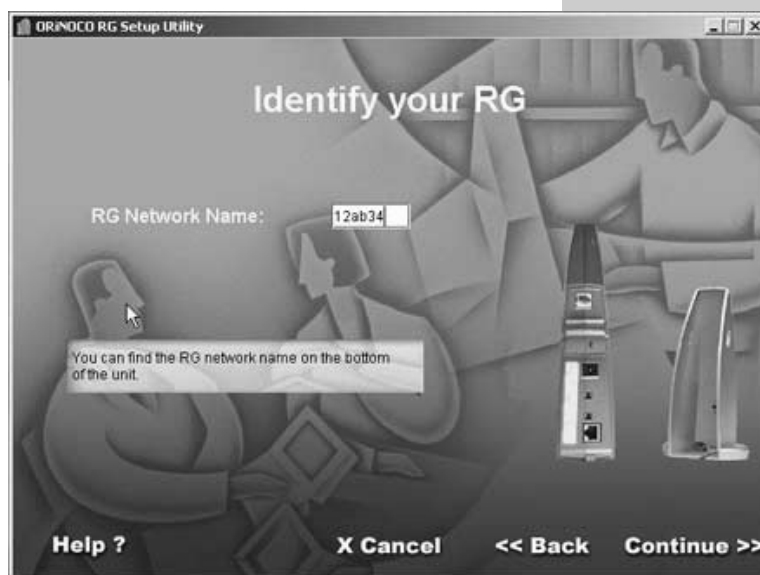
In our case, we installed the setup utility on a machine already on the test network and then started the program.

When the program starts, you are prompted for the network name of the RG-1000, which is a six-character alphanumeric combination printed on a label at the bottom of the device. (See **Figure B**.)

When you click on the Continue prompt, the program will search the network for the RG-1000 with that network name as its default.

During our setup, the software not only found the RG-1000, but it detected older firmware on the device and automatically updated the firmware. This took a matter of moments, and then an access point parameters screen appeared. (See **Figure C**.)

Figure B



You can find the six-character alphanumeric code on the label on the base or back of the unit.

Figure C



Choose the appropriate answers from the drop-down boxes or fill in an IP address in the bottom field.

We changed only the Internet Access Via drop-down box to indicate we were using it over a LAN.

Once the Continue prompt was clicked, a small box opened to allow the setting of the encryption code and password for the access point. The code is the last five characters of

the network name used on the first screen. (See **Figure D**.)

After that, it was a matter of clicking a Finish button in the software, and the RG-1000 was ready to function and wait for a wireless signal from a client.

We had the client software set up on a machine with the USB client, and it was simply a matter of changing its preferences to the Network Name of the RG-1000, setting up the correct encryption code, and checking that it was operating on the same radio frequency as the RG-1000.

Figure D



The Enable Encryption check box is selected by default.

Once those preferences were configured on the client machine, the connection to the RG-1000 was immediate and strong. The client has a signal-measuring capability that is graphically represented on five vertical bars to indicate strength. We received a strong 11-Mbps signal, with four of the five bars filled on the graphic scale.

Games are notorious bandwidth hogs, and the client test machine happened to have a particularly graphic-rich game installed on it. We used the wireless connection to hop through the test network to a game server on the Internet, where we had among the best ping times of any of the contestants playing the game.

The bottom line

If simple setup is the principle criteria in recommending a wireless solution for the small or home office, the new software for setting up the ORiNOCO Residential Gateway makes this product a strong contender.

The biggest difference we could see between the RG-1000 and ORiNOCO's corporate wireless solutions was in the limited number of computers that could be in use with the RG-1000 at one time (10 versus 11 to 25 users for the AP-500 or 26 to 50 users for the AP-1000). ~

ORiNOCO USB client setup makes a turn for the better

Jan 14, 2002

By Mike Walton

To keep from being left behind in the fast-growing wireless market, even the biggest players in this field haven't been able to rest on their laurels. Instead they have been ironing out every wrinkle they can find to help their products succeed in the cut-throat market. Thankfully, ORiNOCO has not

been immune from this pressure and has taken the initiative to work out problems with its product implementation processes. In this review, we'll see if ORiNOCO has fixed the wrinkles in its setup software or if it needs to go back to the ironing board.

Past reviews

In 2001, we examined a number of ORiNOCO wireless products. In our first review, “ORiNOCO’s wireless network: Avoid its sticky setup problems” (page 139), we were disappointed that the company had not done a better job with its setup software. We referred to the setup software as “quirky” and complained that the documentation was too light.

A few months later, when we reviewed an ORiNOCO wireless gateway product, the installation software for the access point side of the equation had been much improved over the client-side installation software. Check out that review, “Installing ORiNOCO wireless gateway is a snap” (page 142), to see how simple the installation was.

Once both of our installations were complete, the ORiNOCO equipment worked as well as advertised. Getting to that point on the client side was the problem.

Improved installation comes to the client side

Given the struggle that we had when previously reviewing the client-side setup software, we were eager to test the new improvements in the client side of ORiNOCO’s wireless LAN installation process.

We already had the ORiNOCO Residential Gateway set up on TechRepublic’s test network for our previous review, so we decided to install the ORiNOCO USB Client (see **Figure A**), a combination that could be used for either home or small office environments.

Our client machine runs Microsoft Windows 2000 Professional operating system, which has USB compatibility built in. Check your operating system to see if it supports USB. If it doesn’t, ORiNOCO offers a PCI card that accepts PCMCIA cards.

It helps to view the installation of the wireless technology on the client side as a four-part process:

1. Install the client manager program.
2. Install the appropriate drivers to connect the wireless client equipment to your operating system.

3. Set the wireless client properties.
4. Set the network properties for your wireless card connection.

Make it all speak the same language

Before connecting your USB or PCI card client hardware to the computer, you will need to insert the installation CD. Once the main menu appears, select Install Software. The software will then take you step-by-step through the process of installing the client manager software.

Next, physically connect the PCI or USB client hardware to your computer and, after you start up again or plug in the USB client, Windows will take you through the Add New Hardware process of loading the appropriate drivers.

This is where we hit the only snag in the process, and that was because we were installing on a Windows 2000 Pro system. When we told the Windows Wizard to get the drivers off the CD, it wanted to install the Windows 98 drivers.

When we forced the Wizard to browse the CD for the Windows 2000 drivers, nothing appeared in the window for those drivers on

Figure A



For this review, we used the ORiNOCO Residential Gateway and USB client, center and right in the photo.

the CD. Stopping the Wizard, we took a look in the CD and found that most drivers were available in the appropriate directories for other operating systems. However, the Windows 2000 driver wasn't in the Windows 2000 directory. Instead, there was a Setup file that allowed us to install the appropriate drivers on our client machine's hard drive. (See **Figure B**.)

After putting all the drivers in a Drivers directory in the ORiNOCO directory where we could easily find them, we unplugged the USB client and then plugged it back in. When the Add New Hardware process started again, we whipped through it without a problem.

After completing the Add New Hardware process, the Add/Edit Configuration Profile

will automatically open to set the properties of your wireless connection. We selected the Residential Gateway access point. For that profile, you type in the Network Name that is printed on the bottom of the access point. After you click OK, you are prompted to restart the computer.

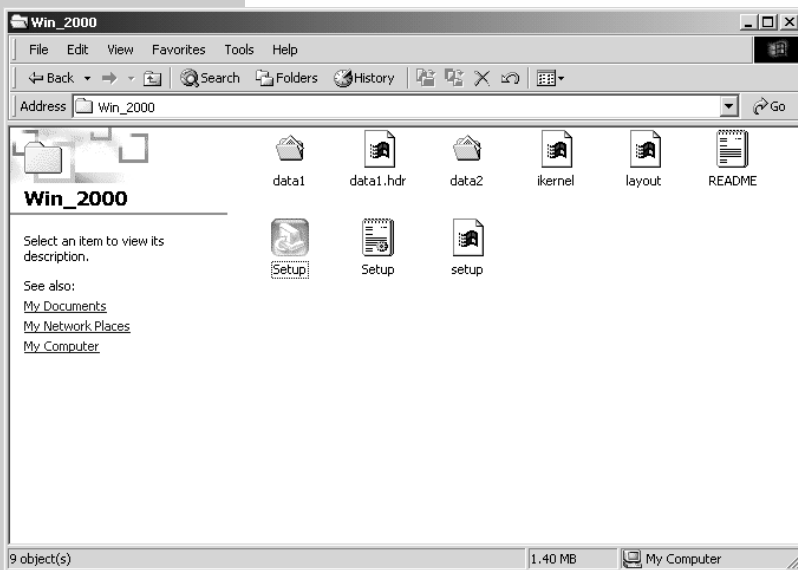
After restarting, when we right-clicked on My Network Places, we had another network connection to configure. For simplicity's sake, we just selected DHCP.

The client manager program places an icon in the system tray that shows the strength of the signal from the access point. Ours was very strong, and when we launched Internet Explorer, we were zooming around the Internet in no time. In our Network Places, other local computers showed up.

The bottom line

If the documentation had told us to use the Setup program in the Drivers folder for Windows 2000 on the ORiNOCO installation CD before plugging in the USB client, the entire process would have taken even less time than it did. As it stands, it was still a fast and simple installation, bringing the client-side implementation up to the high level of ORiNOCO's access-point-side installation. In this case, the fix to the install problem was worth the wait. ~

Figure B



Here is what is in the CD's Windows 2000 directory that did not appear in the Add New Hardware browser.

3Com AirConnect: Wireless for the great wide open

Mar 20, 2001

By James McPherson

Over the next few years, chances are high that you will begin administering wireless networking clients. To help prepare you, we've been running some wireless reviews. 3Com has sent us their 11-Mbps, 802.11b wireless network for review, the AirConnect 11-Mbps PC Card (model 3CRWE737A) and Access Point base station.

The AirConnect products are corporate solutions suitable for industrial settings, complex networks, and corporate campuses spanning acres. The PC Card lists for \$220 and the Access Point retails for \$1,195, but if you shop around, you should be able to find the PC Card for under \$170 and the Access Point for less than \$900. Is it worth your corporate dollar? Read on.

AirConnect technology overview

To help with this review and with questions about encrypted, packetized radio-communication protocols, 3Com provided expert assistance from Paul Keane, a 3Com product engineer. His assistance was greatly appreciated, as it let me go straight to the horse's mouth for clarification.

The AirConnect devices operate on the 2.4-GHz frequency band using the IEEE 802.11b communication standard. This is an internationally accepted standard, granting it a bit more credibility and respect from the business world, not to mention interoperability. Apple originally pushed 802.11b into the limelight when they included 802.11b networking on their laptops under the AirPort brand. Since then, a number of other manufacturers, many of which had 802.11b products prior to Apple's introduction, have begun promoting their products more extensively.

Interference

The 2.4-GHz band is getting very cluttered. Many cordless phones and all microwave ovens operate on that spectrum, as well as

HomeRF and 802.11. Interference is a strong possibility, but 802.11b tries to minimize it by splitting 2.4 GHz into multiple channels and using a variety of transfer speeds.

Roaming

The 802.11b specification is a cellular system, enabling the use of multiple base stations to increase total coverage. Clients automatically negotiate a connection with the nearest compatible base station to maximize connectivity. The system has additional roaming features to allow clients to cross-network subnets and move between routers.

Security

Because different wireless networks could overlap, the base stations have a network identification number (referred to as a wireless LAN, or WLAN, area) to keep clients from wandering into the wrong network. As a result, the security conscious should change their WLAN area from the default number to prevent others from "wandering" into their network. You can also configure specific wireless clients to allow or disallow them access to your network.

Since wireless networking is broadcast to everyone in range, anyone with a receiver has the potential to eavesdrop. To counteract this potential problem, 802.11b includes packet encryption and the option to change security keys—and you should change those keys. There have been recent reports that imply cracking the encryption used by 802.11b will be easier than first suggested. This has yet to be proven, but rotating the encryption keys will help keep your network secure.

Description

The appearance of the AirConnect Access Point is similar to other nonrack-mounted networking hubs (**Figure A**). With dimensions

of 7.5 inches wide and 6 inches deep, the Access Point is slightly larger than most hubs. Its 2-inch high antenna marks the Access Point as a wireless device. Naturally, it includes the standard link, activity, and error lights for debugging.

The Access Point includes hardware to mount the unit on a wall or ceiling, as well as a special power adapter that will let the 10Base-T cable provide power to the unit in locations away from a power supply. (Note: The 10Base-T adapter cannot use a crossover cable to power the Access Point. Doing so will damage the unit.) The power cable is just over 8 feet long and uses the standard three-prong computer power cable, so it can be extended.

The PC Card is similar to other wireless PC card adapters, but it includes a pair of LEDs on the antenna, which is a very welcome addition (**Figure B**).

Other wireless PC cards tend to leave off the LEDs. You can see the PCI version in **Figure C**.

The antenna connects to the PC Card using a pair of metal sockets, which seems fairly solid. The AirConnect seems durable, barring the application of significant force to its antenna.

Abilities

The adapters are, as far as your computer is concerned, just another Ethernet adapter (**Figure D**). Clicking on Properties lets you set your AirConnect Wireless LAN Service Area. Clicking on Advanced allows you to select additional settings for Mobile IP, Encryption, and the WLAN Adapter (**Figure E**).

The AirConnect software provides several welcome testing and monitoring features. With the included network profiler, it is a more-than-complete, if perhaps a too-complicated, solution for all wireless needs. (See the 3Com AirConnect dossier on pages 152-153.)

AirConnect Access Point

The Access Point is very feature-rich. It includes a number of privacy features, such as the ability to use access control lists (ACLs) to lock out particular devices or allow only a select number of devices access. Various types of packet-filtering options are available, including the option to route packets via a modem attached to the serial port.

Despite the extensive number of features, there were several noticeable absences, all of which work to dampen some of my enthusiasm: Without support for PPP over

Figure A




The AirConnect Access Point is about the size of a standard 4-port hub.

Figure B



The sturdy antenna and LEDs make the AirConnect PC Card a cut above.

Figure C



The PCI card is just as sturdy as the PC Card version.

Ethernet (PPPoE), the Access Point is incompatible with many DSL connections. Its inability to act as a DHCP server necessitates a separate server. It's impossible to use the Access Point to perform network address translation (NAT).

Each Access Point supports over 60 clients. These base stations are designed to work in very large arrays, covering multiple acres of floor space to serve hundreds of clients.

If you've read any of my other articles, you're aware of my fondness for security. The default installation uses 40-bit keys—insufficient to deter a truly dedicated data thief but more than sufficient to stop a curious teenager who's figured out how to put his wireless card into promiscuous mode. Fortunately, the U.S. version I tested included 128-bit encryption, which would significantly slow down most data thieves. Nevertheless, because the system uses static keys that have to be manually updated (preventing automated security updates), cracking the security is simply a matter of processing time.

Speaking of security, I don't recommend using the Web interface to administer the

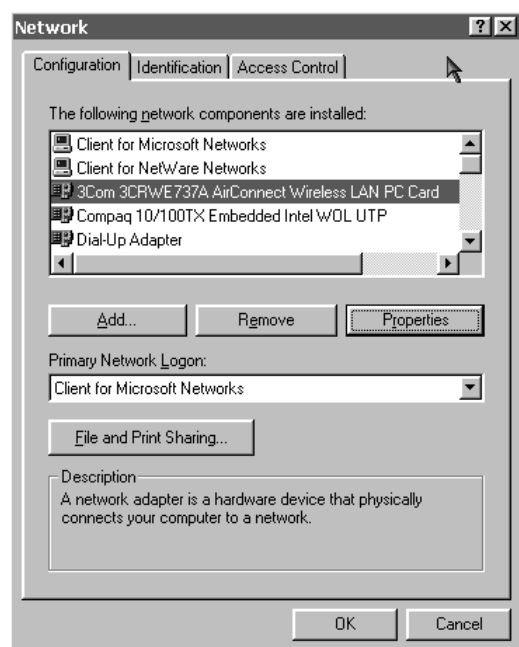
Access Points. The mini Web server does not support SSL encryption, so the administrator password is transmitted in clear text. Anyone with that password can shut down the entire wireless network by doing something as simple as changing the WLAN area or changing the security keys. If someone using the administrator password updates the password and uses the Update All Access Points feature to propagate his changes, you'll be locked out. Resetting the Access Point requires contacting 3Com technical support for instructions—which were intentionally left out of the user manual—as 3Com wisely decided that including them posed a security risk.

Configuration

The Access Point, which must be installed first, can be configured two ways: via serial connection or via a Web interface. The default method is to use a 9-pin serial cable and a terminal program for setup. The installation CD also includes a configuration for HyperTerminal (for the less technically minded).

The serial interface is easy to use, with menus logically laid out and your command

Figure D



This is the Windows 98 Network applet showing the AirConnect PC Card adapter.

Figure E



Advanced Properties settings give you the ability to set your WLAN service area, as well as set encryption, power settings, mobile IP, and others.

sets indicated at the bottom of the screen. In many cases, the screens displayed brief descriptions explaining options.

Using the optional Web interface is only possible if your network is compatible with the default configurations. This interface is much less intimidating and easy to navigate but does not include the option descriptions. An online Help feature is available but requires having the files installed on your PC. I personally prefer to use the included .pdf manuals. Virtually all features of the Access Point can be configured, and you can access all monitored statistics. If you use the Web interface, remember that it does not encrypt the system password.

In all cases, you must have the system password to make changes. Don't lose it. If you do, you'll need to contact 3Com technical support.

Installation of the PC Card was a breeze: Boot the computer, run the installation CD, insert the PC Card when prompted, reboot, and add a dash of the Windows install CD to complete.

Performance

I tested the equipment under different electromagnetic conditions to simulate various office situations. Industrial situations hold a completely different set of conditions, and each location is unique. If you plan on using the AirConnect equipment in an industrial setting, please use the included Site Survey software or hire a Site Survey team.

The stated true performance of the AirConnect system, and all 802.11b devices, is 5 to 6 Mbps at maximum speed (11 Mbps). This is because the system uses a preemptive packet collision avoidance system rather than packet

collision detection employed by wired networks. This preemptive system has an overhead that seems somewhat exorbitant when small numbers of clients are in use. Nevertheless, if you remember that each Access Point is analogous to a 63-port hub, you can see how collision avoidance becomes more necessary as the number of clients increases.

Like a hub, the Access Point shares bandwidth. This means that if you have 60 clients running on a single Access Point, each one would get only about 0.10 Mbps (12 KBps). Of course, this does give each one the equivalent of his or her own ISDN connection, so they won't really notice the limit when accessing the Web or using most types of client-server software. Videoconferencing, however, is probably not an option on a widely used Access Point.

Minimum electronic interference

I placed the adapters and host computers approximately 20 feet from the gateway with two nonstructural walls in between. The only active electronic devices were the local file server (without monitor), a 5-port hub, the Access Point, and a few fluorescent lights. This is as clean a condition as you can expect in a home or office. The results are listed in **Table A**.

These results are a little lower than I expected, but they're not disappointing. Only on network file-sharing operations (like network installs) is a 3.2-Mb connection going to be an inconvenience. Ping times were unaffected under automatic and maximum power consumption modes, and even the minimum power mode provided Ethernet-class latencies.

Table A: *Minimum electronic interference test*

Adapter Setting	Transfer Rate (Mbps)	Transfer Rate (KBps)	Ping Time (Milliseconds)
Maximum power consumption	3.2	400	4
Automatic mode	3.2	400	5
Minimum power consumption	2.0	250	10

Maximum electronic interference

To test a high interference setting, I placed the adapters as before. This time, however, I cranked up the juice. Static generators consisted of a 21-inch monitor located about 18 inches from the gateway, a desktop PC, the file server and hub, a 32-inch TV, the computer and adapter, and the coup de grace, a 300-watt microwave about 15 feet down the hall operating on High. (I thought about scuffing my feet on the carpet and arcing into the doorknob but figured that was going too far.) I expect that this is more indicative of the working conditions with which this device should expect to contend. As expected, bandwidth and latency suffered some. **Table B** shows the results.

Network performance dropped 25 percent with active interference. Packet loss occurred with some regularity, and the wireless link dropped from 11 Mbps to 5.5 Mbps more than once. Of course, not many places will have this type of concerted interference to deal with, but don't be surprised if people near the kitchen complain from time to time.

I was initially disappointed at the idea of only half the feasible bandwidth until I realized that the worst I could throw at it would slow it down to only half speed, which is still a throughput high enough to completely monopolize a T-1 internet connection. At that moment, I realized how much this technology could spoil a person.

Range test

The typical maximum office range of 25 meters (75 feet) for an 11-Mbps connection might be a tad optimistic. My office is not that large (8 meters x 12 meters), yet I managed to

find a place where the AirConnect could not connect at faster than 5.5 Mbps. In other words, your mileage may vary and you would be wise to expect 802.11b devices to have a maximum 11-Mbps range of 12 meters.

My open line-of-sight connections were limited to about 125 feet (40 meters). The connection remained 11 Mbps. I increased the stress by walking around a nearby building, and the connection dropped down immediately to 5.5 Mbps. It didn't stay there long, as the numerous intervening brick walls, interior walls, and electrical conduits caused the connection to wobble between 2.2 and 1 Mbps. Packet loss was an issue but was noticeable only to local servers; Internet sites performed as usual.

Issues

I have high expectations for most products, even more so for an established manufacturer like 3Com. No performance problems cropped up with this one; in fact, the devices worked flawlessly for the two weeks I tested them. And I did use them—constantly. My laptop saw more use than ever before as I could wander about at will, always connected. It was addictive.

I have a huge issue with the lack of security on Access Point passwords, though. At no time is the password encrypted; it is always broadcast in clear text. Anyone with that password can reprogram the Access Point. Using the option to propagate changes to other Access Points, a hacker could hijack the entire network.

I also question 3Com's decision to have all encryption disabled by default. While I understand that some users may be confused by the

Table B: Maximum electronic interference test			
Adapter Setting	Transfer Rate (Mbps)	Transfer Rate (KBps)	Ping Time (Milliseconds)
Maximum power consumption	2.4	300	4
Automatic mode	2.4	300	5
Minimum power consumption	1.7	225	10

security settings, the settings are quite simplistic. The only potential pitfall is not using the right encryption key, but this is just a matter of clicking the right button. Only people who change the encryption key will have to make any significant effort.

Also, some network statistics aren't reported correctly. The encryption statistics report no encrypted packets sent, but it also reports encrypted packets received. According to 3Com, this is an error in the UI.

Final grade

The Access Point is a complex component, on par with the 3Com SuperStack series of Hubs. The price is steep, but when you compare one

\$900 Access Point (capable of handling 63 clients) to three \$200, 24-port SuperStack hubs and the associated wiring, you can see the value—especially in wide open areas that would require long cable runs.

The 3Com devices are significantly more expensive than their competition. Online, you can find 802.11b PC cards and Access Points for \$115 and \$300, respectively, saving \$50 per card and over \$500 for the base station. Whether their devices are up to the same standard is impossible to tell as of yet, but 3Com has put a lot of effort into their products, and the 3Com name has a lot going for it.

The performance was quite acceptable, and the network-side options were excellent. The

3COM AIRCONNECT DOSSIER

EQUIPMENT TESTED

- ▶ AirConnect Wireless LAN PC Card, model 3CRW737A Rev. B
- ▶ AirConnect Access Point
- ▶ 3Com AirConnect software version 1.5 (Site Survey, Status Monitor, drivers)

SPECIFICATIONS

Cellular system allows roaming between Access Points and across network segments

RANGE (PER CELL)

- ▶ Maximum clear open air: 300 meters
- ▶ Typical Office Max (11 Mbps): 25 meters
- ▶ Typical Office Max (1 Mbps): 90 meters
- ▶ Maximum clients per cell: 63
- ▶ Packet encryption with configurable keys: 40-bit international, 128-bit United States

MULTIPLE BANDWIDTH MODES

- ▶ Original 802.11: 2.2 Mbps (0.26 MBps), 1.1 Mbps (0.14 MBps)
- ▶ Updated 802.11b: 11 Mbps (1.4 MBps), 5.5 Mbps (0.7 MBps)

ACCESS POINT FEATURES

- ▶ DHCP client and server operation
- ▶ SNMP-enabled

- ▶ Access control lists for user administration
- ▶ Packet-filtering services include forwarding of trapped packets via modem
- ▶ Network statistics
- ▶ Accessible via serial port, telnet, and Web
- ▶ Dual antennas providing signal diversity to help isolate signals
- ▶ Mobile IP services to allow roaming between Access Points on different routers
- ▶ Network statistics available through Web interface
- ▶ Configurations can be propagated to other Access Points

PC CARD FEATURES

Six power settings: Continuously Active Mode (CAM) and five Power Saving Polling (PSP) modes. PSP modes cause the client to power down and wait a number of network cycles before requesting data. CAM clients request data every cycle.

SOFTWARE

AIRCONNECT STATUS MONITOR

As the most significant part of the included software, this versatile application provides quick and easy access to almost any

features were thorough for those seeking to build a wireless campus. Mobile IPs enable incredibly large (physically or IP-wise) networks to work across routers. ACLs, packet filtering, and the ability to propagate settings across the network are requisite for anyone running a complex network.

Client software was more than complete. The network profiler is a vital tool but is more complex than I think necessary. The necessity to have telephony and dial-up networking upgrades on a Windows 95 system seemed extreme. Including the Site Survey application and the diagnostic features of the Status Monitor, however, are valuable bonuses.

Documentation was poorer than I've come to expect from 3Com. I found a few noncritical errors between the Quick Start guide and the user manual. The lack of documentation for the Windows CE platform was disheartening.

I give the Access Point a B. It's a workhorse that lacks a little polish. The PC card and client software receive an A-; the cards worked wonderfully, but I was troubled by the complexity of the Mobile Connection Manager (MCM) software. This software is necessary for Windows 9x machines whose users have to switch connections. With so many current and future Windows 9x systems on laptops, the complexity of MCM will be a real difficulty for end

information or setting pertinent to the operation of the adapter.

- ▶ Bandwidth mode
- ▶ Ping times and lost packets to a particular host
- ▶ Number of local Access Points
- ▶ Signal strength to Access Points
- ▶ Signal noise of each Access Point
- ▶ Communication channel in use
- ▶ Data transfer statistics
- ▶ Power mode used by the PC Card

MOBILE CONNECTION MANAGER

An undeniably powerful, but also overly complicated, network profile utility was included that enables Windows 9x/Me portables to readily migrate from location to location without requiring network reconfigurations. The system supports all networking devices, including modems and remote printers. Unfortunately, this requires installing telephony and dial-up networking upgrades to Windows 95 clients to support the variety of services. This software is useful to anyone who travels to a large number of network sites. It includes the ability to import and export profiles, which can help

standardize system configurations. Future versions need to be easier to use.

SITE SURVEY SOFTWARE

This package focuses on placement of Access Point units to provide consistent coverage of a facility. It consists of the Status Monitor software with a more flexible interface and logging to evaluate the connection quality between multiple Access Points and record the test client's location. Site Survey is a useful tool to help identify dead zones within a coverage area. You can contract a Site Survey team from 3Com services.

DRIVERS

Drivers were included for Windows CE 2.0, 95, 98, 2000, and NT. No documentation was included for the Windows CE drivers. According to 3Com, v1.5 drivers are compatible with Windows CE 2.11. New v2.0 drivers, which were due to be released February 2001, will support Windows CE 3.0.

NEW FEATURES


Support for peer-to-peer connections between individual AirConnect PC and PCI adapters will be included in the version 2.0 drivers, which were due out February 2001.

users. The Proxim HomeRF I reviewed included an incredibly easy-to-use variant of the MCM that worked perfectly. Therefore, I know it is possible to build and is relatively inexpensive. Nevertheless, MCM software is targeted at mobile users who visit a large number of diverse network environments and need full functionality.

The 3Com AirConnect product line is excellent for corporate wireless programs spanning large areas or multiple buildings, such as a corporate campus. The feature set is appropriate for organizations with an existing network and IT staff that can ensure that the advanced features provided will be taken advantage of.

For small or midsize offices looking for integrated solutions, check out 3Com's other

802.11b base station, the Home Gateway. The Home Gateway supports NAT, acts as a DHCP server, and supports PPPoE and VPNs. The upgraded version of the client software, which will include peer-to-peer networking using wireless adapters only, will be more suited to those looking for an integrated solution. (As of this writing, the upgrade, scheduled to be released February 2001, has not yet been released.)

AirConnect is 3Com's 18-wheeler of networking; it rattles a bit and may not be the easiest to drive, but it will get a whole lot of data where it's going. 

A review of 3Com's HomeConnect Home Wireless Gateway

May 29, 2001

By James McPherson

We have the sequel to the 3Com Access Point base station; it is called the HomeConnect Home Wireless Gateway. Unlike the campus-focused Access Point, the HomeConnect is intended to be used as a standalone, wireless gateway and is equipped with the features you'd expect: The DHCP, NAT, firewall, PPPoE, and a three-port, integrated 10/100 Mb switch make this device perfect for the SOHO or limited-wireless-needs corporate network audience. The HomeConnect is priced competitively, listing for \$299 when I searched for it on Price-watch.com.

Technology

The HomeConnect Wireless Gateway device operates on the 2.4-GHz frequency band using the IEEE 802.11b communication standard. The 802.11b standard is internationally accepted as a standard that covers many products from a wide variety of manufacturers, all of which guarantee to provide basic compatibility. Coincidentally, 2.4-GHz is getting to be a very cluttered band. Many cordless phones and all microwave ovens operate on this spectrum. Interference is a possibility, but 802.11b tries to minimize it by using multiple channels in

the 2.4-GHz band, as well as by having a variety of transfer speeds.

While the 802.11b is a cellular system that supports the use of multiple base stations to increase the total coverage zone, the HomeConnect is not equipped with this feature. It retains the network identification numbers (referred to as a WLAN area) that keep a client from wandering into the wrong network. As a result, the security conscious should be sure to change the WLAN area from the default to prevent others from wandering into their network.

SECURITY TIP

Since wireless networking is broadcast to everyone in range, anyone with a receiver has the potential to eavesdrop. To counteract the security concern, 802.11b includes packet encryption and the option to change the keys in use. You should change the keys. There have been recent reports that imply cracking the encryption used by 802.11b will be easier than originally suggested. This has yet to be proven, but rotating the encryption keys will help keep your network secure.

Physical description

The HomeConnect Gateway has a very unobtrusive flat black design; no protruding antennas, wings, vanes, or unusual post-modern art additions. The face is equipped with the standard link, activity, and error lights for debugging. Three RJ-45 ports for the internal network, a single RJ-45 uplink port, and the reset button adorn the back.

Specifications and concerns

Wireless support is limited to 35 clients, half the Access Point's 60 clients. The three internal network ports offset that and provide enough capacity for your average wired small office network, especially if you want to cascade switches and hubs off it. However, you might want to hold off before planning to use the Gateway as your main network component.

Network features include a DHCP server, network address translation, a firewall, and client service filtering. These are consumer-grade SOHO features not suited for the enterprise, and here's why:

First, the DHCP server is locked in to the 192.168.2.x Class C IP block, putting a total limit of 253 clients and factoring out the gateway's use of 192.168.2.1. You can't manage the IP leases manually or force the Gateway to use a single IP for a specific MAC address. While suitable for most SOHO applications, it does limit the reliability of the client filtering, as client filtering is based on IP addresses.

That forced IP address creates another potential snag. While the network address translation (NAT) feature via DHCP is very nice, it isn't really an option. Oh, you can use a different server to manage your NAT needs, but with that internal address always forced to 192.168.2.1, you really can't put the Gateway into transparent networking mode.

Client filtering is intended to keep your kids or employees from getting into things they shouldn't. Filtered protocols consist of Web (HTTP), mail (POP3 and SMTP), news (NNTP), FTP, and telnet and can have start and stop times, along with day-of-week restrictions. These settings are easily understood but provide little flexibility. Any client that configures another IP or configures a proxy on a nonfiltered port can get around the client filter. Since you have to configure each IP separately, you can't make blanket settings.

The firewall has an internal log feature and is targeted to block nine common types of network attacks: five denial of service (DoS) and four intrusion techniques. The DoSs blocked are the SYN land attack, SYN flood, Smurf ICMP broadcast echoes, Snork UDP packet routing, and the classic oversized ping. The firewall also claims to defend against UDP port scans, zero length IP packets, TCP null scans, and IP spoofing, which is admittedly more of a deception than an intrusion.

I tried several of the ping attacks and all were readily deflected. Since I don't often attack other netizens, I don't have the other attacks handy, but I'm fairly certain the listed

tools will be blocked. I would be more concerned by the attacks that aren't listed. However, it is more likely that any DoS attack will take down the Gateway instead of your PCs.

I also scanned the firewall to check the logging functions and was somewhat disappointed. Only my SYN scan was detected; the stealth TCP and UDP scans went undetected, as did the use of IP spoofing. A sense of false security is a dangerous thing.

As far as the physical security of your communications, the Gateway uses the default 802.11b 40-bit security keys. This level of security is more than sufficient to stop a curious teenager who figured out how to put their wireless card into promiscuous mode but really won't slow down a determined data thief. However, even using the default encryption keys would at least stop random passersby with 802.11b devices from borrowing your bandwidth without some modicum of work.

Installation

Setting up the HomeConnect Gateway is like placing any other simple switch or hub in a network. Locate it near a power outlet, connect RJ-45 cables to the ports, make sure to get your upstream source cable in the clearly marked and separate uplink port, and the wired connections are good to go. Wireless configuration is done through a simple Web interface. The limited number of options makes it difficult to mess things up and the layout is very functional. Documentation is thorough and is perfect for the target audience.

Performance

The wired performance of the Gateway was up to par, recognizing 10 Mb and 100 Mb connections and transmitting data at about 98 percent of the rated speeds. Not bad, but 3Com's been selling 10/100 Mb switches for years, and anything less would be disappointing. Besides, if you're buying this product, the switch is a bonus on top of the wireless functionality, which is where we will spend our time.

Wireless tests were conducted under different electromagnetic conditions to simulate various office situations. The HomeConnect is not intended for industrial situations. If you need wireless equipment in an industrial set-

ting, you should look at the AirConnect Access Point.

The stated performance of the HomeConnect system, and all 802.11b devices, is 5 to 6 Mbps at maximum speed. This is because the system uses a preemptive packet collision avoidance system rather than the normal packet collision detection system employed by hard-wired networks. This preemptive system has an overhead cost that seems somewhat exorbitant when small numbers of clients are in use. However, if you remember that the HomeConnect's wireless client support is analogous to a 35-port hub, you can see how necessary collision avoidance is.

Like a hub, the Gateway shares bandwidth on the wireless segment among all clients. This means that if you had 30 clients and the Gateway was working at its maximum theoretical speed of 5 Mbps, each one would get only about 0.15 Mbps (20 KBps). Of course, this would give each one the equivalent of an ISDN connection. However, that is assuming the Gateway works at the theoretical maximum. At the risk of spoiling the next section, I must admit that's not a good assumption.

Testing conditions

Minimum internal electronic interference

The only active electronic devices were the local file server (without monitor), the Gateway, and a few florescent lights. No computers were attached to the internal switch. This is as clean a condition as I could expect to appear in any home or office.

Maximum electronic interference

Interference has been introduced in the form of a 21-inch monitor located about 18 inches from the Gateway, a desktop PC, the file server and hub, a 32-inch TV, the computer, Gateway, and the coup de grace, a 300-watt microwave about 10 feet down the hall operating on high. I thought about scuffing my feet on the carpet and arcing static discharges into the doorknob but figured that was going too far. I expect that this is more indicative of the indoor working conditions this device should expect to contend with.

Range and obstructions

In addition to the introduced electronic interference, I tested the device at a variety of ranges. The number and type of obstructions are also noted. Tests over 30 feet were made with the Gateway within six feet of a window and the client outside. Realize that the effects of the electronic interference are more noticeable when in close proximity. At longer range you will be dealing more with what your client has to deal with rather than what is near the Gateway. **Table A** shows the results of my testing.

As you can see, network performance under best-case conditions was disappointing. Throughput was never able to exceed half the Gateway's maximum functional bandwidth (5 to 6 Mbps). This is fine for typical consumer DSL service, but it is a little slim for file sharing within the office or taking advantage of larger connections.

Range didn't really challenge the Gateway, as the performance was adequate around 100 feet. When in immediate proximity, the signal quality dropped significantly but not enough to really interfere with normal operation. However, when used at any range worth justifying a wireless connection, the performance degrades to nearly unusable levels with the introduction of interference. Thus, you are wise to look at the layout of your office and make sure you

won't be using the Gateway anywhere near your break room. Regardless, the Gateway isn't going to let you surf the Net in the parking lot of your building or in your backyard unless you happen to park in direct sight of it or leave the window open.

Final grade

While the Gateway is no replacement for an actual security policy implemented by people who understand how to secure and maintain a network, it does provide an additional layer of security that will help protect end users. It should block the common attacks and scans that plague many cable modem and DSL networks without affecting your internal network adversely.

Wireless communication has an inherent security risk that comes from transmitting a signal willy-nilly into space. However, the use of 40-bit signal encryption, rather than the more powerful 128-bit encryption, is acceptable in a product of this class. I feel it was less acceptable to have all encryption disabled by default. The only potential pitfall is using the right encryption key, but since default settings are, well, default, it's just a matter of clicking the right button. Only people who change the encryption key will have to make any significant effort, and those people should

Table A

Range (in feet)	Obstructions	Interference	Listed Mbps	Transfer Rate (Mbps)	Transfer Rate (KBps)	Ping time (in milliseconds)
						Min/Max/Avg/Lost
5	None	Minimum	11	2	250	4 / 4 / 6 / 0
5	None	Maximum	5.5	1	125	4 / 14 / 7 / 0
30	Two interior walls	Minimum	11	2	250	4 / 4 / 7 / 0
30	None	Maximum	1	6 KB	.5	4 / 17 / 7 / 0
100	None	Minimum	5.5	1.4	175	4 / 9 / 4 / 0
100	None	Maximum	5.5	1.4	120	4 / 12 / 6 / 0
150	None	Minimum	2	.5	64	5 / 55 / 9 / 1
150	None	Maximum	1	0.123	16	5 / 153 / 22 / 5
100	Two exterior walls	Minimum	1	Unsustained	Unsustained	5 / 100 / 22 / 10

realize the work entailed, which, for the record, consists of typing the same 10-character string into the Gateway and into each of your clients.

The Gateway is a nice piece of equipment as long as it is used when and where intended. It may seem like I'm overstressing this; however, if you expect too much from the Gateway, you will be disappointed. It is targeted for small sites with little to no current network infrastructure that need a flexible single solution that doesn't require much maintenance. It does all that quite admirably, although I think the performance and signal strength were a little weaker than I would like, especially after seeing the Access Point.

The HomeConnect Gateway provides a cost-effective 802.11b base station suitable for SOHOs that complements the AirConnect Access Point in 3Com's product lineup. The price is comparable to other feature-rich, consumer-grade base stations once you factor in the integrated switch and the advantage of 3Com's considerable reputation. So if you need a standalone wireless network solution that can handle all your basic needs, the Gateway is the toy for you. If you plan on integrating it into a larger network, you should pay more attention to its older sibling and know enough to stay away from SOHO products. ~


3Com Home Wireless Gateway

Sep 10, 2001

By Allen Fear, ZDNet

3Com's \$399 Home Wireless Gateway measures just 7 x 8.5 x 2 inches, and it's shaped sort of like a cigar box. Yet this compact, unassuming device functions as the digital nerve center of your home network. It communicates with notebook and desktop

computers over radio waves, allowing them to share a high-speed Internet connection, as well as printers and files, within a 300-foot range at speeds up to 11mbps. (Note: While the Home Wireless Gateway lets you share an Internet connection among both PCs and Macs, you'll need to run Windows NT Server's Services for Macintosh or a comparable service if you want



EDITORS' RATING

7.0

Value	7
Features	8
Design	7
Installation	6

PROS

- Easy to install
- Wi-Fi compliant
- Integrated firewall
- Supports cable

CONS

- No telephone jack.

Specs

Model first available	January 1, 2001
Linux compatible	No
Mac compatible	Yes
PC compatible	Yes
Standards supported	802.11b
Interface	Ethernet

to share files or printers between the two operating systems.)

In addition to the Home Wireless Gateway, remember that each computer on your network must have a wireless Ethernet adapter. Since the Gateway supports the 802.11b standard and is Wi-Fi certified, it is operable with Wi-Fi cards from 3Com and other vendors.

The Home Wireless Gateway is easy to set up and manage. To set up the device, you simply plug in the power supply and connect the included Ethernet cable to your DSL, cable, or ISDN modem connection. Next, install a network adapter in each computer you want to network, and configure the TCP/IP settings for communication with the Home Wireless Gateway. There's no software to install; the device includes an integrated configuration tool that you access over a standard Web browser. Just enter the provided IP address into the address bar of your browser. When the Setup program appears, go to the Gateway Setup Wizard and follow the onscreen instructions. The included Installation Map and User Guide also provide step-by-step instructions.

Aside from supplying wireless connectivity, the Home Wireless Gateway also has three wired 10/100 Ethernet jacks for faster data-transfer speeds, in case you want to swap large video files. To connect to a wired Ethernet port, your computer must have a network interface card (NIC) installed.


Our experiences testing the Home Wireless Gateway reflect the possibilities—and limitations—you may encounter in your own home. For example, 3Com claims you can wirelessly connect up to 35 computers using the Home Wireless Gateway. That big a cluster, however, could easily slow the network to a crawl, because all machines connecting to an 802.11b Wi-Fi network have to share available bandwidth. Based on the results of our tests, the Home Wireless Gateway is probably best suited for Wi-Fi-compliant networks of seven or fewer clients. When you factor in the device's three Ethernet ports, you have ten nodes total, more than enough for most home-networking environments.

3Com also claims that you can roam wirelessly anywhere within a 300-foot range of the Home Wireless Gateway, but in our ZDNet Labs' tests, a range of 200 to 250 feet was more realistic. Also, the gateway automatically reduces transmission speeds to 5 MBps, 2 MBps, and finally 1 MBps, depending on the quality of the signal, so the farther you are from the device, the slower the connection speed will be. Ultimately, the range will depend on the acoustics of your home.

The Home Wireless Gateway includes an integrated firewall that uses network address translation (NAT) and an array of defense techniques to protect you against many of the most common attack methods. This is not an industrial-strength firewall, but it should be adequate to protect your network from eavesdroppers. The unit also safeguards data with 40-bit encryption, and it offers VPN (virtual private network) pass-through support for the most common security standards and protocols, including L2TP, PPTP, and IPSec.

We had only one complaint about the Home Wireless Gateway. Because the unit is targeted at a broadband audience, it doesn't include a modem or an RJ-11 phone jack. If you want to share a wired or wireless dial-up connection, check out ORiNOCO's RG-1000 Residential Gateway.

Kudos to 3Com for backing the Home Wireless Gateway with a lengthy, five-year warranty. Toll-free phone support is fairly convenient for the home user. It's available Monday through Friday, 6 A.M. to 9 P.M. and Saturdays from 9 A.M. to 3 P.M. Also, the 3Com Web site offers a searchable knowledge base, manuals, FAQs, and downloadable firmware upgrades.

3Com's Home Wireless Gateway is an ideal way to share a broadband connection. This Wi-Fi-compliant device is easy to set up and configure, and it requires no software installation. It securely protects your data and even lets you access your office network over a VPN. For the price, it offers a good solution for a busy home's Internet-access and other sharing needs. 

Connect wires and wireless with the Linksys Ethernet Bridge

Jun 5, 2003

By William C. Schmied

Networks have been around now for more than 20 years. But in the small office/home office (SOHO) community, I've seen that many networks today aren't exactly networked. Here's an example: I have a client who has a small Windows 2000 domain running with one DC and about a half dozen clients. These machines are all relatively close to each other—close enough to run Cat 5e cable and tie them together using a 100-Mbps switch. These computers have good network connectivity with each other and all is well. But add to this mix two Macintosh OS 10.2 computers and two Windows XP computers that have been placed 100 feet away. Now we have a quandary. How will we get these four additional computers on the wired LAN affordably? Enter the Linksys WET11 Wireless Ethernet Bridge.

Buying the bridge: Cost justification

Now I know what you are thinking. Why not just give each of these computers a wireless

network connection? After all, AirPort cards for the Macintosh can be found in the \$80 to \$100 range, and several vendors—including Microsoft, Linksys, NetGear, and SMC—offer low-cost USB network adapters for Windows PCs that are in the \$50 to \$80 range.

But after you've spent \$260 to \$360 or so on wireless network adapters, you will still need to get a good access point. Again, you have many choices, so let's just keep it simple and use the Linksys WAP11 as an example. You can grab one of these for about \$80 to \$110 just about anywhere. So our total cost is now in the range of \$340 to \$470 dollars to bring these four computers onto the network. That's pretty hefty and will continue to grow as we add more clients in the future. Fortunately, there is an alternative.

The price tag on the WET11, as of this writing, was between \$100 and \$130. Add that to the \$80 to \$110 cost of the WAP11, plus the EZXS88W 8-port 10/100 switch at about \$45 to \$55, and you have a total solution in the \$225 to \$295 range. So you've already saved money—always a good thing. And you're supporting only two new network devices (the WAP11 access point and the WET11 bridge; the switch requires no support) instead of five new network devices (the WAP11 and four wireless network adapters). Any time you can minimize the number of devices you have to support and configure, the better off you are. So it certainly looks like implementing the WET11 wireless bridge with the WAP11 access point and eight-port 10/100 switch is going to solve the problem and make my job easier. With all the pieces in place, let's set it up and see how it goes.

Building the bridge

Out of the box, the WET11 comes with everything you need to get going. It includes the bridge itself, a power cable, a standard Ethernet cable, a removable antenna you can change out

Figure A



The box includes everything shown here.

if desired, a CD-ROM containing the configuring utility software, and a quick start installation guide that should be all you need to get the WET11 installed and operating. A detailed user's guide is available in PDF on the CD-ROM, as well as from the Linksys Web site.

Figure A shows the contents of the box.

What's my mode?

The first thing you need to do is to figure out how you will be connecting the WET11. Will you be connecting it directly to a computer or other computing device (such as an Xbox or PlayStation 2) or will you be connecting it to a hub or switch? You'll need to configure the WET11 to operate properly (in cross-over or straight-through mode) depending on how it will be used in your network. As you can see in **Figure B**, the markings on the back of the WET11 are simple enough. Move the switch to X for cross-over (to connect directly to an Ethernet device, such as a computer or PlayStation 2) or to II for straight-through to connect the device to a hub or switch. In our case, we will be using a switch, so we need to move the switch to the II position.

Connecting the pieces

After you've configured the bridge to operate in straight-through mode, you'll need to connect its various parts. This is also a good time to get your clients connected to the switch and get the switch powered up. You'll want to make sure that you connect the Ethernet cable from the WET11 to the uplink port on your switch. If you don't, clients won't be able to connect through the bridge to the access point.


Beam me up, Scotty!

Well, we're finally ready to get down to the business of using the WET11 bridge. What you do next depends on whether you've already gotten your access point configured and operating yet. **Figure C** shows the pertinent configuration information you will need to know about the WAP11. The WAP11 includes a fairly easy-to-use setup wizard if you need to get it set up.

When configuring any new wireless networking device, you'll want to keep in mind:

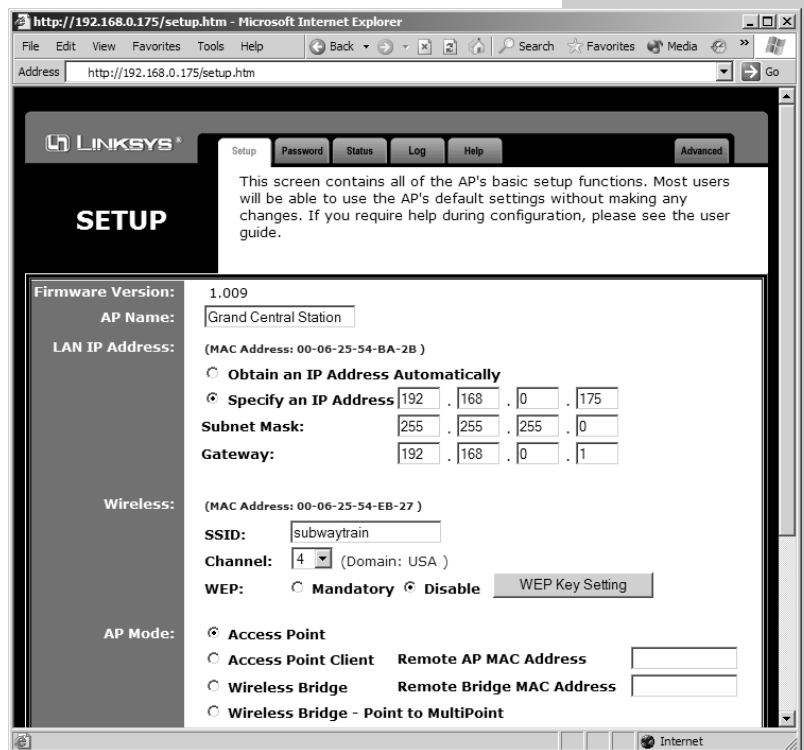
- **The IP address you will be assigning to the access point.** It's a good idea to stati-

Figure B



Don't forget to select the bridge's mode of operation here first.

Figure C



Here's some information you may need to be aware of.

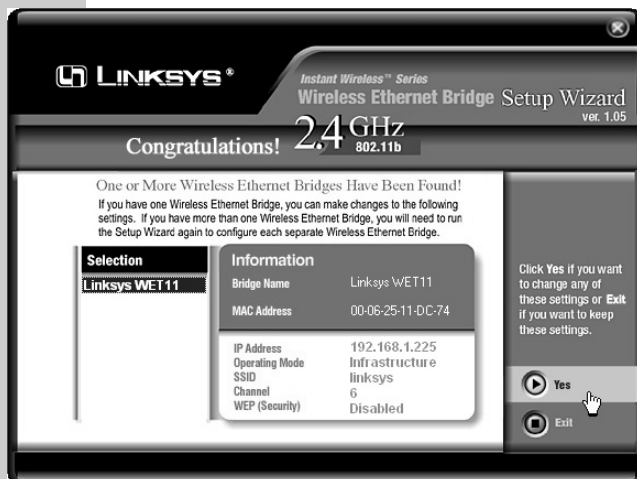
- call assign this just the same as you would with any other infrastructure device.
- **The default gateway IP address.**
- **The subnet mask value.**
- **The network SSID you will be using for your wireless network.** This is essentially equivalent to the Windows workgroup name but should not be easy to guess.
- **The channel that you want your wireless devices to operate on.** Which one you choose in a small environment with only one access point does not matter; if you have multiple access points, you will need to

Figure D



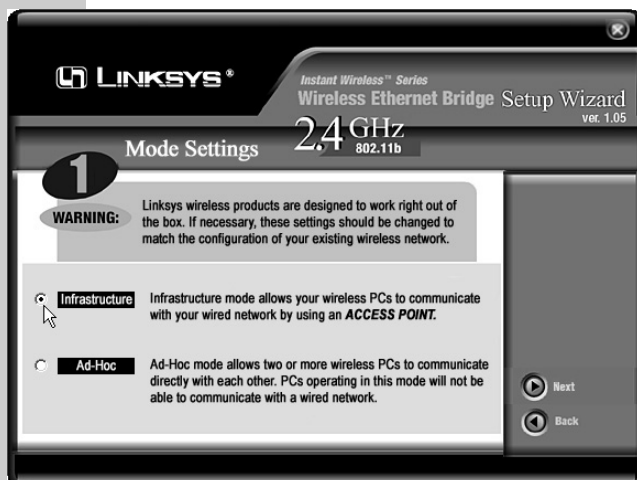
When you start the WET11 Setup Wizard, you'll see this screen.

Figure E



These settings will need to be changed.

Figure F



Choose Ad Hoc mode when you first see this screen.

be concerned with channel overlap and should consider using channels 1, 6, and 11 only.

► **The WEP status (either off or on) and the keys in use.**

For the WAP11 specifically, you will need to ensure that it is configured for Access Point mode. The WET11 bridge will simply be acting as a normal wireless network client, as far as the WAP11 is concerned.

After connecting everything, you'll need to place the included CD-ROM into one of your computers so that you can begin the bridge setup. The best choice is a computer that has a wired connection to the bridge through the switch. If the setup routine doesn't auto-run, start it by double-clicking the Setup.exe file on the CD. You'll be greeted with the Setup Wizard screen, shown in **Figure D**.

Clicking Setup will prompt the Setup Wizard to scan for the WET11. After the scan is complete, you will be presented with the results page, which should look similar to that shown in **Figure E**.

You will need to run through the rest of the Setup Wizard to change the settings to match those of your wireless network before you'll be able to use the WET11. A problem that I've noted with all Linksys wireless network hardware is that it does not automatically change the channel to the one in use when placed in Infrastructure mode (as any wireless infrastructure device should). As a result, you will need to run around in a circle for a bit to get things straightened out. From the Mode Settings screen, shown in **Figure F**, you'll need to select Ad-Hoc mode and click Next.

In the Basic Settings screen, shown in **Figure G**, you can now configure the correct SSID, channel, and device name settings. Click Back to return to the page shown in Figure F, select Infrastructure mode this time, and click Next. You will be brought back to the Basic Settings screen again, but this time, you'll notice that the channel selection is grayed out—the reason why we went in a circle.

Finally, you will be presented with a summary screen showing your configured settings, such as the one shown in **Figure H**. This would be a good time to record this information for

future reference. Click Yes to save your settings. You will be prompted to unplug the power cord from the bridge for a few seconds and then plug it back in to enable the settings to take effect. I found in my installation that the settings took effect instantly and clients could surf the Web immediately.

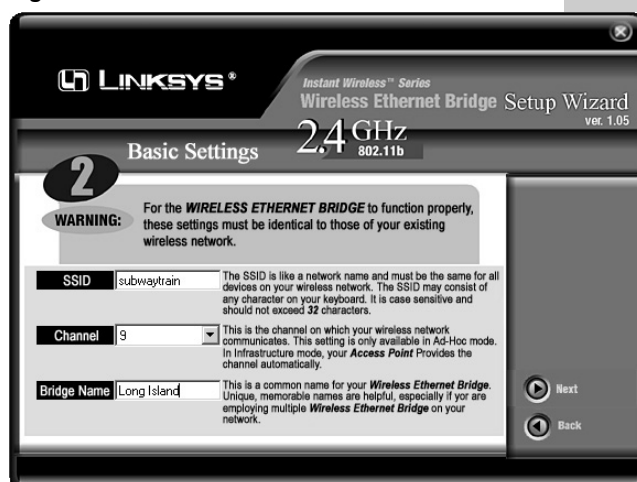
Bridging the gap

Overall, the WET11 bridge is a solid solution for a SOHO environment that needs to connect computers to the wired network without the hassle of running a new length of Ethernet cabling. You can also easily connect your Xbox or PlayStation 2 to the Internet using the Bridge. After all, not too many people have Ethernet near their PlayStation 2. The setup is pretty straightforward and should be complete in about 20 minutes or less.

As with all wireless network connections, you should seriously consider implementing WEP protection on the transmissions. Although it is true that WEP has been cracked, to not use it is really just asking for someone to penetrate your internal network.

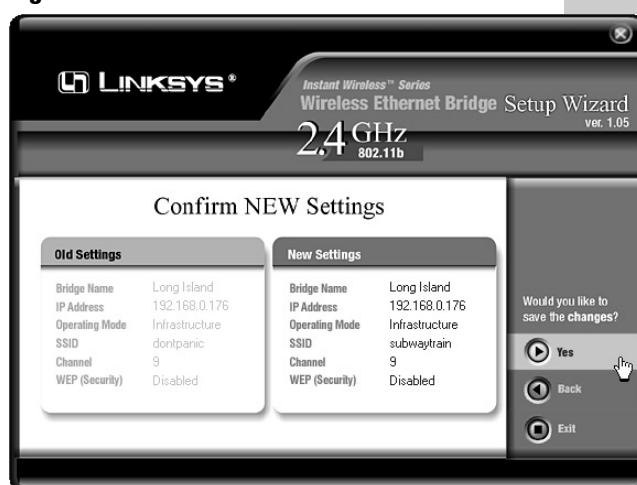
My only real complaints about the WET11 are the standard issues that Linksys products have with selecting the channel number and the fact that it does not have the same standard form factor as most of the other network products Linksys manufactures. Because of this, it does not stack well and may tend to get knocked around or slide around. You will want to place it in a location where it is not subject to being bumped. And as soon as you can, be sure to go back into the browser-based admin panel by entering the IP address you assigned the WET11 and change the password to something a bit more secure than the default of *admin*.

Figure G




This is your only chance to configure the channel during the Setup Wizard.

Figure H



The last step is to verify the settings and save the configuration.

The WET11 Wireless Bridge presents a good solution to a common problem many people may have. Best of all, it does it at an affordable price. 

Linksys EtherFast wireless AP and cable/DSL router with 4-port

Jan 8, 2002

By Mark Henricks, ZDNet

If you want to share a broadband cable or DSL connection across a home- or small-office network with cabled Ethernet and wireless 802.11b segments, the Linksys EtherFast wireless router does a solid job. It combines a four-port hub with an 802.11b wireless access point, a cable/DSL router, and a firewall. Although the package has some shortcomings in documentation and security, it still provides convenience and speed, replacing several pieces of equipment that would cost much more.

A multifunction network device

Linksys includes everything you need to set up your home network in one easy-to-install package. The \$229 EtherFast wireless router basically combines a wireless 802.11b access point with Linksys' hot-selling four-port cable/DSL router. The unobtrusive design has four sturdy legs supporting a rounded, rectangular black-and-purple case. Should you wish to connect two or more routers via the uplink port, recesses on top simplify stacking. The wireless access point acts as a DHCP server and assigns IP addresses to PCs on the network. It also supports WEP encryption and claims a top operating range of 300 feet (91 meters) indoors and 1,500 feet (457 meters) outdoors.


The Ethernet switch operates at 10 Mbps or 100 Mbps and has four LAN ports, plus a WAN jack for the modem. Front LEDs indicate power status as well as broadband, wireless, and cable activity. A Reset button on the back lets you restore the router to its default factory settings. In addition to the EtherFast wireless router, the kit includes a power adapter, a CD-ROM with software and documentation, and a printed user guide.

Simply plug it all together

Installing the EtherFast wireless router was relatively painless. We came across some confusing sections in the one-page quick-installation guide, but fortunately, the excellent and well-detailed 60-page user guide answered all of our questions. To get started, connect the router to its power source, the cable/DSL modem to the WAN port on the back of the router, and the installation PC to one of the LAN ports. Finally, set the TCP/IP settings for the installation PC's network card to obtain an IP address automatically, and then reboot. Once you've successfully installed the router on your network, you can configure it using your Web browser. Simply type the provided IP address and password to launch the router's setup page.

Watch your network take off

Performance was great when it came to throughput, notching 92.5 Mbps on the Ethernet connection and 4.2 Mbps wirelessly in CNET Labs' tests. Wi-Fi compatibility was also seamless. The EtherFast wireless router worked as well with an ORiNOCO 802.11b wireless PC Card as it did with Linksys' own cards. Range was about as good as we've seen: with 75 feet and several walls separating the EtherFast wireless router and the wireless clients, signal strength fell marginally, but most messages passed at the top 11-Mbps rate, and dropped information was minimal.



EDITORS' RATING

7.3

Performance	7
Design	8
Installation	7
Features	7

PROS

- Great phone support
- Excellent performance
- Smooth installation

CONS


- Lacks key security feature
- Confusing documentation

Specs

Maximum theoretical throughput	10 Mbps, 100 Mbps, 11 Mbps
Connectivity	Four 10/100Base-T Fast Ethernet LAN, one 10Base-T WAN, one shared uplink port
Warranty on parts/labor	One year
Device type	Router
Protocol(s)	Ethernet
Network type	10Base-T Ethernet, 100Base Ethernet, 802.11b
Compatible operating systems	Windows 95, 98, NT, 2000, or Me
Supports DHCP	Yes
Model first available	May 7, 2001

As a router, the Linksys offers support for IPSec pass-through, PPTP (point-to-point tunneling protocol), PPPoE (point-to-point protocol over Ethernet), and DMZ (demilitarized zone) mapping. You can set up the router to filter Internet access (handy for family home networks), allow remote administration, keep a log of all Internet sites visited, and more. Unfortunately, unlike the D-Link DI-714 wireless broadband router with four-port switch, the Linksys EtherFast doesn't include stateful packet inspection among its security features, which would have provided an added level of security to the existing NAT and TCP/IP inspection.

The Linksys EtherFast comes with a one-year warranty. Although the warranty may be standard, the router's phone support is above average. Toll-free phone support is available 24/7, excluding major holidays, for the life of the product. The Linksys Web site offers firmware updates, a searchable knowledge base, user guides, and FAQs.

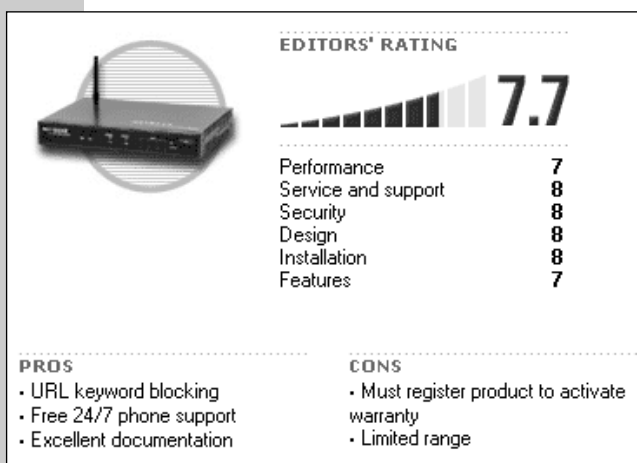
The Linksys EtherFast wireless router could have done a better job with its setup documentation and offered more complete security features. Nevertheless, for the price, it still is a good solution for SOHO users looking to combine wireless and cabled network segments with broadband Internet access. 

NetGear MR314 cable/DSL wireless router

Apr 24, 2002

By Allen Fear, ZDNet

If you're worried about what your kids may be viewing on the Web (and who isn't these days?), then the NetGear MR314 cable/DSL wireless router may be your family's new best friend. With its Web-filtering capabilities, the MR314 lets you restrict access to Internet content you deem objectionable or inappropriate. However, this device offers more than just peace of mind to nervous parents. It also features an easy-to-use, Web-based configuration utility; an integrated 802.11b access point; a four-port Ethernet switch; inbound firewall protection; and WEP encryption, all of which make it a good fit for any home looking to share a wireless broadband connection.



Specs

Maximum theoretical throughput 100 Mbps/11 Mbps

Connectivity Wireless, cable

Warranty on parts/labor Five years

Device type Router

Protocol(s) Ethernet, Fast Ethernet, IEEE 802.11b

Compatible operating systems Windows, Mac

Model first available June 12, 2001

Thirty minutes or less

Using the installation guide, we got the \$248 MR314 up and running well within the estimated 30-minute setup time. To get started, connect your wired computers to any of the four Ethernet ports located on the back of the router. Next, use the included CAT-5 cable to connect your cable/DSL modem to the router's Internet port. In addition to the router, you must also purchase an adapter (such as the \$140 NetGear MA101 802.11b wireless USB adapter) for each wireless desktop or notebook you want to connect.

To configure the router, open your PC's browser and type in the provided IP address. When the main menu appears, select the Wizard Setup and follow the onscreen instructions. (To help you configure your Internet settings, NetGear provides a handy ISP guide that tells you exactly what information you need.) If you get stuck along the way, an exhaustive Reference Guide on the included Resource CD covers configuration settings and has a troubleshooting section, a glossary, and a brief overview of networking basics.

Router and nanny in one

The MR314 has one interesting feature that sets it apart from other home routers. From the configuration utility's Advanced menu, you can block access to certain sites based on specific URLs, keywords within the URLs, or time of day. For example, you can ward off sites with the words "playboy" or "sex" in the address. You can even have the MR314 send an e-mail to you if an attempt is made to connect to a site with your listed keywords.


In addition to basic network address translation (NAT, which hides your computers' IP addresses), the MR314 offers other security features you'll want to activate. From the configuration utility's Advanced menu, you can enable 64- or 128-bit WEP encryption, change the administrator's password, or restrict network access using MAC address filtering. The

MR314 supports dynamic DNS and port forwarding for those who need more sophisticated capabilities, and you can also set up one PC as a DMZ.

Great throughput, limited reach

The MR314 turned in excellent performance on our labs' tests. Its wireless throughput of 4.9 Mbps and Ethernet throughput of 88.3 Mbps tied it with the SMC Barricade wireless broadband router. The NetGear's range was slightly weaker than that of other routers in its class; however, in an indoor environment with some intervening walls, we achieved stable connections at distances of up to 60 feet.

The MR314 comes with a lengthy five-year warranty and toll-free, 24/7 phone support. But there's a catch; to activate the warranty, you must register your product within 30 days. Otherwise, support lasts only 90 days. The company's Web site offers FAQs, downloads, and contact information.

If you want to monitor what your kids see on the Net, then the MR314 is a good choice. It also offers an easy-to-use, Web-based configuration utility and plenty of performance. And while we could do without the product registration hoopla, we like the five-year warranty. 

HP wireless gateway hn200w

Apr 8, 2002

By Eric Knorr, ZDNet


The HP wireless gateway hn200w has the coolest appearance of any such device on the market. Even the setup program looks great. With its friendly automatic setup and use of nontechnical terms, this device is clearly geared toward network novices and home users.

In addition to an 802.11b radio, you get four 10/100-Mbps Ethernet ports (more than most units), an uplink port for attaching another hub, and an Ethernet cable. It's too bad this unit disappointed us by refusing to work with Windows XP and by burying features in its configuration software. The one-year warranty was also a letdown.


A software maze

Unlike most wireless gateways, the \$220 HP requires that you install software to get started. However, incompatibility caused by the app kept the unit from working with Windows XP,

and HP's support was unable to fully identify the problem (we even downloaded a beta firmware upgrade to no avail). HP expects to resolve all XP incompatibility issues by June 2002, but for now, we used Windows 2000 to test the gateway instead. The unit comes with an illustrated Quick Start Guide, which covers



EDITORS' RATING

 **6.4**

Performance	7
Service and support	7
Security	6
Design	8
Installation	5
Features	6

PROS

- Includes network cable
- Excellent 802.11b throughput

CONS

- Incompatible with Windows XP
- One-year warranty
- Confusing software

basic hardware installation, but for more detailed information, you'll need to refer to the comprehensive user guide located on the CD-ROM.

The gateway's setup process began relatively smoothly. When you insert the CD and start the installation wizard, the software automatically grabs your ISP's settings instead of asking you to enter them manually—a nice touch. Once you've installed the Gateway Control Panel, however, you're faced with a tabbed, quirky interface that uses nonstandard language (unprotected in place of DMZ, for example) and hides basic features. While novices may appreciate the nontechnical wording, we found it ultimately more confusing. Anyone familiar with routers will head straight for the Expert Interface tab, which provides browser access to settings. There, you'll find such essentials as MAC address cloning, without which some ISPs might stop you cold.

Accent on access


Beyond basic network address translation (NAT), which hides your computers' IP addresses, HP sacrifices security for accessibility. For example, the manual recommends against turning on either 64- or 128-bit WEP encryption (if you can find it) because of the resulting performance hit. Additionally, there's no function that alerts you to unauthorized attempts to hack your network. At least you can go through the Expert Interface and disallow individual MAC addresses or grant access to only those with MAC addresses on your list.

The most unusual twist is the elaborate set of options for controlling access rights on your network. The Gateway Control Panel makes it easy to disable Internet access for any computer on your LAN. But it gets more granular than that—you can even block individual applications (such as multiplayer games) from accessing the Internet. Even better is the Gateway Control Panel's ability to make specific applications available over the Internet with a few clicks, without demanding that you mess around with port forwarding.

Performance makes the grade

In our labs' tests, the HP tied with the Belkin wireless cable/DSL gateway router and the Siemens SpeedStream, all three of which achieved 802.11b throughput rates of 4.9 Mbps, more than enough for home use. A Windows XP system ran the tests for Belkin's and Siemens's units, however, so we couldn't be sure that running Windows 2000 with the HP (as we were forced to do) affected performance. The unit's range and ability to penetrate walls were right up there with those of the best 802.11b gateways. And the \$120 HP 11-Mbps wireless LAN PC Card we used in our tests was a snap to set up. HP also sells a convenient \$130 wireless USB network adapter for desktop use.

The meager service and support for HP's gateway is another reason we can't recommend it highly. Its mere one-year warranty is outstripped by the longer guarantees offered by competitors. Free phone support is available 24/7 during the warranty period, but you'll have to pay toll charges. The Web site is helpful, though, with FAQs, manuals, software updates, and other useful information.

In the end, the HP wireless gateway hn200w amounts to a failed attempt to create an easy home-networking appliance. It works conceptually, offering a sweet design and some automatic configuration, but the execution could use some work. If HP fixed the XP problem, redesigned the software, and provided a longer warranty, the company could have a hot little gateway on its hands. 

Specs

Maximum theoretical throughput	100 Mbps/11 Mbps
Connectivity	Wireless, cable
Warranty on parts/labor	One year
Device type	Wireless access point
Protocol(s)	Ethernet, Fast Ethernet, IEEE 802.11b
Compatible operating systems	Windows 98 and above
Model first available	January 2, 2002

Intel AnyPoint wireless gateway

Apr 8, 2002

By Eric Knorr, ZDNet

Resembling a broad, flat, beige mushroom, the Intel wireless gateway does a decent job of providing 802.11b wireless access for home users mostly interested in surfing the Web.

On the downside, the setup is a bit quirky, only one Ethernet port is provided, and Intel offers limited phone support. The upside: Its speed is decent, the documentation is pretty complete, and—if you need the option—the Intel doubles as a wireless access point. The three-year warranty is enticing, too.

It's in the air

The \$219 Intel comes with a good, illustrated quick-start guide. You plug the unit into your DSL or cable modem and, through the unit's one LAN port, hook up your computer via a crossover Ethernet cable that Intel graciously supplies. From the outset, Intel gives you the option to set up the gateway as a wireless access point on an existing network—a waste of the router capability but a convenient option if you need it.

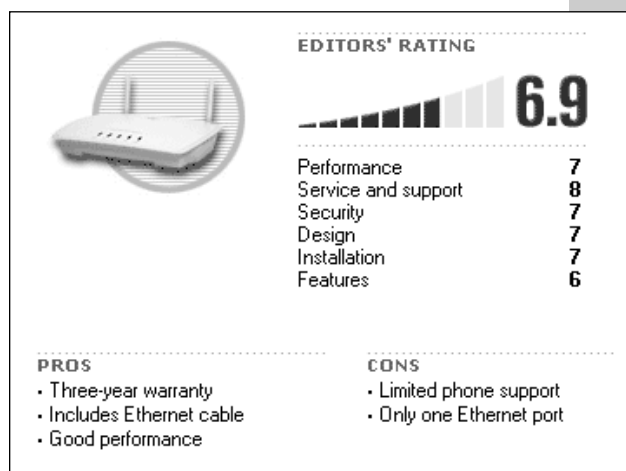
The Intel boasts a serial port that might lead you to believe that it can connect to a dial-up modem as a backup—but no such luck. The manual says that the port is for “advanced users to view or change the gateway's settings using Telnet or HyperTerminal interfaces” instead of the configuration firmware. No explanation is offered to describe how you might do this.

To communicate with the gateway, you'll need a wireless network adapter for each computer you plan to connect. Intel, like most wireless gateway vendors, sells and recommends its own adapters. We used the \$99 AnyPoint Wireless II Network PC Card and the \$109 AnyPoint Wireless II Network USB model. For those on a budget, the company also offers “entry level” home networking adapters that top out at 1.6 Mbps (compared to 11 Mbps in the regular models) and cost only \$30 a pop. If all you want to do is share Internet access, the low-priced models

should be sufficient for most broadband connections.

Staying safe and sound

To configure the gateway, insert the CD-ROM and follow the onscreen instructions. The setup wizard works pretty well, asking you for ISP settings, prompting you for a network ID code, and so on. The full documentation on CD-ROM is somewhat disorganized, but it covers the bases pretty well, including a glossary and a troubleshooting guide. During setup, the configuration utility asks you to create your own password. You get basic NAT security, of course, but the wizard also recommends WEP encryption, with detailed instructions for entering 64- and 128-bit keys. You



Specs

Maximum theoretical throughput	100 Mbps/11 Mbps
Connectivity	Wireless, cable
Warranty on parts/labor	Three years
Device type	Wireless access point
Protocol(s)	Ethernet, Fast Ethernet, IEEE 802.11b
Compatible operating systems	Windows 95 and above
Model first available	October 1, 2001


can either set up one PC as a DMZ or slog through port-forwarding settings to open specific apps to Internet access. In addition, you can screen out certain MAC addresses or open your network to only the MAC addresses you specify.

Running with the pack

The Intel delivered perfectly respectable 802.11b throughput on our labs' tests. At 4.6 Mbps, it was bit slower than the Belkin wireless cable/DSL gateway router and the HP wireless gateway hn200w, but it was still fast enough for everyday home use. (The unit has only one Ethernet port, so testing wired Ethernet speed was pointless.) As expected, the router won't achieve its maximum 300-foot range unless no walls stand in the way, though we saw no real difference in

range or penetration among the three units tested.

The gateway comes with an above-average, three-year warranty. But it's too bad that free phone support lasts for only 90 days after purchase; after that, you pay \$2.50 per minute or \$15 per incident. Intel's Web site offers software updates, FAQs, and e-mail support.

Considering Intel's reputation as a tech giant, the wireless gateway somewhat disappointed us. The speed is fine, and we like the three-year warranty. However, installation and configuration are a bit awkward, and no extra features really distinguish the package. Instead, we recommend competitors such as the Siemens SpeedStream wireless DSL/cable router, which offers more for the money. 

SMC Barricade wireless broadband router

Apr 24, 2002

By Allen Fear, ZDNet

Considering its name, you might expect the Barricade to come with retractable metal grating encased in barbed wire. Instead, SMC's wireless broadband router is a sleek, gray box about the size of a big slice of deep-dish pizza. But despite its modest appearance, it offers more connection ports than any other home router we've tested, with the exception of the Siemens SpeedStream Wireless DSL/cable router. It also offers excellent performance and an impressive set of security features to protect your network from the most common hacker attacks.

Ports galore

The \$178 Barricade offers an impressive array of ports. It includes three 10/100-Mbps Ethernet ports and a 10-Mbps WAN port for connecting your cable/DSL modem. The Barricade also has an RS-232 serial port for

connecting to an ISDN terminal adapter or a POTS analog modem; the latter is a handy backup if your broadband service becomes temporarily unavailable. The Barricade also has a parallel port, but because new printers typically connect via USB, most home users will have difficulty finding a use for it.

The Barricade offers some other conveniences as well. Two omnidirectional antennas extend from either side on the back to provide better range, and a Reset button next to the printer port sends the router back to its factory default settings. While the Barricade does not include brackets for wall or ceiling mounting, it does come with a CAT-5 Ethernet cable.

Simple administration

Setting up the Barricade is simple. The Quick Installation Guide contains detailed and illustrated configuration instructions for PCs and

Macs, and a more comprehensive user guide in PDF format can be found on the included CD-ROM. To configure the router, open your Web browser, type in the provided IP address, and click the Setup button from the main screen. The automated setup wizard leads you step by step through the rest of the process. To connect wireless computers, you'll need to purchase an adapter (such as the \$60.95 SMC EZ Connect 11-Mbps wireless USB adapter) for each machine.


Good security

You can also make more sophisticated networking settings from the Advanced menu tab. For example, you can configure the Barricade to function as a virtual server for services you would like to set up behind the router's NAT-based firewall. A simple check box tells the router to discard pings from the WAN side, which helps conceal your router on the Internet. You can also run multiuser applications behind the firewall by opening public ports or assigning a particular machine to run without firewall protection within a DMZ. Other tabs let you update the firmware, reset defaults, or check the security log, where you can view any illegal attempts to access your network. The Barricade's firewall can also block common hacker attacks, including IP spoofing, land attack, ping of death, smurf attack, and snork attack.

The Barricade performed admirably in our labs' tests. With 4.9 Mbps of wireless throughput and 88.3 Mbps of Ethernet throughput, it matched the NetGear MR314 cable/DSL wireless router. In informal range tests, the Barricade delivered better and more consistent signal strength than the MR314 when connecting through walls, but by only a few feet.

Lasts a lifetime?

SMC's complex warranty and support policies make the user work a bit to get the best deal. The Barricade comes with a standard 90-day warranty, but you can upgrade to a limited lifetime warranty if you register your product



EDITORS' RATING

7.9

Performance	7
Service and support	8
Security	8
Design	8
Installation	8
Features	8

PROS

- Serial modem port
- Two omnidirectional antennas
- Print server
- Solid firewall protection

CONS

- No ceiling or wall mounting hardware
- Quirky warranty

Specs

Maximum theoretical throughput	100 Mbps/11 Mbps
Connectivity	Wireless, cable
Warranty on parts/labor	Limited lifetime
Device type	Router
Protocol(s)	Ethernet, Fast Ethernet, IEEE 802.11b
Compatible operating systems	Windows, Mac
Model first available	April 1, 2001

within 30 days. Limited lifetime means SMC will support the product for up to one year past the date the company decides to discontinue the product. After that, warranty repair or replacement is considered on a case-by-case basis. Toll-free phone support, however, is available 24/7 for as long as you own the product. The SMC Web site also offers drivers, FAQs, and e-mail support.

If you need a wireless router for your home or small office, SMC delivers an attractive package at a reasonable price. The Barricade offers excellent performance, a plethora of ports, and good security to boot. But we do find the carrot-and-stick support policies rather harsh. ☹

SMC EZ Connect 802.11a wireless access point

Apr 29, 2002

By Mark Henricks, CNET

CNET Rating: 7 out of 10

With the SMC EZ Connect 802.11a wireless access point, you no longer need to sacrifice financial security or networking performance when you give up wires. This device's modest price and good performance make it suitable for network gaming and streaming video over short-range, wireless networks. However, first-time networkers may find the EZ Connect's sometimes balky configuration and limited reach troublesome.

A good value

At \$356, the EZ Connect 802.11a is a good solution at a competitive price. The access point's gently curved, plastic casing sports two omnidirectional antennas, three front-panel LED lights, and connections for Ethernet and power. A Reset button, which returns the access point to its factory settings if the Web-based configuration utility fails, nestles between the two jacks on the back. The package also contains a power adapter; a helpful, 40-page manual; and a CD-ROM containing the same manual in electronic form, drivers, and the EZ Connect 802.11a Configuration Utility, which you use to access the Web-based configuration page. To test the unit, we used a laptop outfitted with SMC's \$143 EZ Connect 802.11a wireless Cardbus adapter.

Tricky configuration

As its name implies, installing the EZ Connect 802.11a is remarkably easy; however, configuring the unit is a bit trickier. First you install the EZ Connect 802.11a Configuration Utility on the PC you plan to connect; use your own Ethernet cable to connect your PC to the access point (SMC does not include one); then run the utility. The application should find the access point automatically; unfortunately, it didn't. It worked only after tech support rec-

ommended we set the wired PC to an IP address similar to the access point's default address. Apparently, this is a common problem with setting up the access point. Inexplicably, SMC fails to address it in the manual. You can also manually get to the Web-based configuration screen by typing the access point's provided IP address into your PC's Web browser.

Once the setup wizard was up and running, it was easy to use. SMC walks you through specifying the SSID, enabling turbo mode, and implementing 64-, 128-, or 152-bit WEP security. The advanced setup screen lets you set the access point to work as a DHCP client or server. And you can modify settings for synchronizing with other access points or set data-packet sizes. If your network suffers from signal interference, shrinking the packet sizes increases network reliability—but reduces its efficiency. The status screen displays more than two dozen useful bits of information, including MAC address, WEP status, mode (turbo or regular), and signal strength, in easy-to-read tables.

Solid performance

Like other 802.11a access points, the EZ Connect 802.11a operates in the 5-GHz band, free from cordless-phone and other device interference that can plague 802.11b networks. The EZ Connect also has a top speed of 54 Mbps, or nearly five times the 802.11b benchmark. In CNET Labs' tests, however, it produced just less than 21 Mbps of throughput—average among 802.11a access points. Proprietary turbo mode theoretically boosts speed to 72 Mbps, but in tests, it actually yielded just 25.6 Mbps, which is middle-of-the-road compared to other manufacturers' turbo modes. Like all 802.11a turbo implementations, SMC's turbo mode won't work with other manufacturers' equipment.

Range was also a challenge for the EZ Connect 802.11a. In our workout, it fell short of

Product specification

General

Device type	Wireless access point
Compatibility	PC
Form factor	External
Software included	Drivers & Utilities

Networking

Connectivity technology	Wireless, cable
Networking compliant	IEEE 802.3-LAN, IEEE 802.3U-LAN, IEEE 802.11a-LAN standards
Data link protocol	Ethernet, Fast Ethernet, IEEE 802.11a
Status indicators	Link activity, power, link OK
Features	128-bit WEP, 64-bit WEP, 152-bit WEP

Expansion/connectivity

Port(s) total (free)/	1 (1) x network Ethernet 10Base-T/100Base-TX / RJ-45 female - 1
Connector type	1 (1) x network Radio-Ethernet

Physical characteristics

Width	7.4 Inches
Depth	1.0 Inch
Height	5.2 Inches
Weight	0.2 Pounds

Power

Compliant standards	UL
Power supply included	Power adapter - external
Voltage required	AC 100/240 V (47/63 Hz)
Voltage provided	3.3 V

System Requirements


Minimum operating system	Microsoft Windows 98, Microsoft Windows 2000, Microsoft Windows NT, Microsoft Windows Millennium Edition, Microsoft Windows XP
--------------------------	--

its stated range of 1,650 feet outdoors and 165 feet indoors. Speed dropped off rapidly once a few walls intervened. Separated by 60 feet and a floor, the EZ Connect saw its transmission rate drop to 6 Mbps, followed by a lost connection. If range is a major concern, you may want to look elsewhere.

Read the warranty

SMC's service and support policies for the EZ Connect 802.11a are generous, provided you read the fine print. The access point comes with a standard 90-day warranty, but you can upgrade to a limited lifetime warranty if you register your product within 30 days. Limited

lifetime means SMC will support the product for up to one year past the date the company decides to discontinue it. After that, warranty repair or replacement is considered on a case-by-case basis. Other support basics are more straightforward; the company provides free, 24/7 phone support, and the Web site offers drivers, FAQs, and e-mail support.

SMC's EZ Connect 802.11a has its shortcomings: a shaky installation and slightly tricky support policies, most notably. But if you are willing to tolerate these rough spots, you'll enjoy this access point's easy-to-use, Web-based interface and solid throughput—not to mention its affordable price. 

Quickly add wireless ports with SMC's EZ Connect wireless access point

Jan 6, 2003

By John Sheesley

With all the talk about security and deployment problems with wireless access points (WAPs), the thought of adding a WAP to your network may send your blood pressure rising. WAPs are supposed to make it easier for users to get their work done, but they invariably add to the network administrator's workload. Fortunately, deploying WAPs needn't be stressful. SMC's EZ Connect wireless access point lets you quickly set up and secure a WAP on your network.

WHAT WAP?

For the purposes of this article, I'm going to discuss the SMC EZ Connect wireless access point, model number SMC2655W. This is a basic 802.11b WAP that can connect users using any 802.11b compliant device. To test this WAP, I used ViewSonic's V1000 Tablet PC with its integrated Intel 802.11b networking card.

Setting up the WAP

SMC has stripped the EZ Connect right down to the basics. There are no firewalls to worry about, nor any other switches or wired

ports. This lack of additional features make the WAP easy to set up and administer, but it can raise the final cost if you need some of these features because you'll have to buy them separately.

WAPs don't get much easier to set up than SMC's EZ Connect. All you have to do is plug the WAP in and connect it to your network. Like a workstation, the WAP can either connect directly to a wired switch or a patch panel, so long as the panel is patched to a hub or switch.

According to SMC's specifications, the EZ Connect can connect users up to a distance of 1,800 ft. As with most things, your mileage will vary. The actual distance and speed you'll get will depend on how you deploy the WAP.

During my testing, I couldn't effectively connect the Tablet PC to the WAP at distances over 100 ft, but that's because the WAP was set up on a desk, not placed in a high location like SMC recommends. In addition, the offices at my workplace use metal studs, which can block radio signals. Therefore, any distance problems weren't the unit's fault.

In addition to placing the WAP in a high location, SMC recommends that you orient the

dual antennas for maximum coverage. One antenna should be vertical, while the other should be laid horizontally.

Using the EZ Connect Wireless AP Manager

Once you've plugged the WAP in, you can configure it by using SMC's EZ Connect Wireless AP Manager. You'll find the utility on the floppy disk that comes with the WAP. To install the Manager, open a command prompt on your administration workstation and run Setup from the Utility directory on the floppy.

Running Setup is just as easy as physically hooking up the WAP. You won't find any surprises; Setup runs just like every other Windows installation wizard you've ever run. Follow the steps in the wizard, clicking Next and making your choices along the way.

After you've installed the Manager, you're ready to configure the WAP. Click Start | Programs | EZ Connect Wireless AP Manager | EZ Connect Wireless AP Manager. When the Manager starts, it will start scanning your network for your EZ Connect WAP.

Don't panic if the Manager doesn't see your WAP initially. It's also not a problem if you see the WAP in the Manager but get an error when you first connect to it. As you can see in the IP Address field, the WAP starts off with a default TCP/IP address of 192.168.0.254. If this addressing scheme conflicts with your network's addressing scheme, your administration workstation can't connect to the WAP.

To fix this problem, just temporarily readdress the administration workstation's TCP/IP address to match the WAP's scheme. For example, you may want to change the administration workstation's address to 192.168.0.1. When you do, the Manager will successfully connect to the WAP.

To connect to the WAP, select it from the Manager and click the Connect button. You'll then see the Input Password screen. Type the default password, *MiniAP*, in the password field. When the connection is successful, you'll see Connected in the State column.

The first thing you should do is to change the default password (MiniAP). To do so, select Change Password from the Command

menu. When the Password Configuration window appears, enter the new password and click OK.

The next thing you should do is to change the WAP's TCP/IP addressing scheme to match your networks. To do so, select Change AP from the Command menu. You'll then see the AP Setting menu.

The WAP can use DHCP if you have a DHCP server on your network. You can enable DHCP by selecting Enabled from the DHCP Client drop-down list box. Even if the list box shows Enabled, your WAP may still have the default 192.168.0.254 address. Select Disabled and set an address that matches your network's scheme. Then click Save.

When you do, the WAP will reset with the new address. You'll need to exit the Manager and reconfigure your administration workstation back to its old TCP/IP address. You can then reconnect to the WAP using the Manager and complete your configuration. Don't forget to enter your new password rather than the default one of MiniAP.

Securing the WAP

When you again connect to the WAP, you can configure it for security. Begin by setting an SSID for your WAP. All of the WAPs in your network should share the same SSID. Likewise, you'll need to set this SSID on any devices that will connect to the WAP. To do so, enter a value in the SSID field. This can be any alphanumeric value. Along with setting an SSID, you should remove the check from the Accept Any SSID field. This will prevent the WAP from accepting devices that don't have the proper SSID.

Next, set a unique channel for the WAP. In the United States, the WAP can run off any of 11 different channels. To avoid interference, try to set one that's different from any channels that already exist in your area.

An important part of security in a wireless network is WEP (Wired Equivalent Privacy). Even though WEP gets a lot of criticism for not being truly secure, it's certainly better than nothing. To set WEP, click Encryption.

When the Encryption Setting screen appears, select 128 Bits from the Encryption

(WEP) drop-down list box. However, you should do this only if your wireless devices can support 128-bit encryption, so check your devices first. You may be forced to use 64-bit encryption or, worse, no encryption at all.

To set the password, select Create With Passphrase and enter a passphrase in the Passphrase field. Manager will hash the passphrase and create a WEP key. You can then reuse this passphrase on wireless devices to rehash the matching key without having to write down the 26-character key.

If you have existing WAPs in your organization, you can manually enter the key by selecting Manual Entry and carefully typing in the key in the fields provided. As you may have noted when you used the passphrase, you'll be temporarily able to see the resulting 26-character key in the Manual Entry fields. You may want to write this key down. Some wireless

devices won't properly hash the passphrase to create a matching key, so you'll need to manually enter it on those devices. This was a particular problem on the ViewSonic Tablet PC.

After you've set WEP encryption and entered the other security information, you can close Manager. The WAP will reset, and you'll be able to connect to it from your wireless devices.

Less wire, less hassle

Setting up a wireless access point for your network needn't be a stressful event. Manufacturers are making it easy to deploy WAPs with units like the EZ Connect wireless access point. Just make sure you take the time to properly configure and secure your access point, and you'll have your users wandering around wirelessly with ease. ~

Untether your network with SMC's wireless adapter

Oct 21, 2002

By Ray Geroski

One of the newest products in the SMC EZ Connect line is the SMC2664W 2.4-GHz USB wireless adapter. In conjunction with a wireless access point, you can use this adapter to quickly and easily connect any Windows PC with a USB port to the network. Because it connects via USB, the SMC2664W is a great solution for both laptop and desktop users who need a wireless connection. It's easy to install and, for the most part, works as advertised. If you're looking for an alternative to a PC Card adapter, the SMC2664W is a solid option.

Parts and installation

The SMC2664W package includes the following:

- ▶ The wireless adapter with internal antenna
- ▶ Driver and utility disks
- ▶ 6-ft. USB cable
- ▶ Fastening clip
- ▶ Velcro swatches
- ▶ Two small magnets

You'll notice that the drivers and product utility software come on disk, not CD. This could have been a problem for me because the laptop on which I was testing the device didn't have a floppy drive. My workaround was to copy the contents of the disks onto a USB storage device. Given my access to the USB storage device, the disks presented only a

minor inconvenience, but some users might find this a bigger obstacle if they, too, don't have floppy drives. My advice to SMC: Put the software on a CD.

Of course, updated drivers and the product manual are both available for download on SMC's Web site (<http://www.smc.com/index.cfm?sec=Products&pg=Product-Details&prod=251&site=c>).

Some assembly required

Depending on how you plan to use the adapter, you may have some small parts to assemble. A clip included in the package snaps on to the back of the adapter, allowing you to attach it to your laptop monitor. The manual says you can even clip the adapter to your belt, although I'm not really sure why you'd want to unless you were going to walk around with your laptop. Velcro pads and magnets are also included for mounting in various locations.

It's pretty simple to install the adapter—you just attach the USB cable and plug it in to your PC. Windows will then detect the device and prompt you to install the driver. Depending on which OS you're using, you'll then navigate to one of two folders on the driver disk: One folder contains the drivers for Windows 98, Windows 2000, and Windows Me; the other contains the Windows XP drivers.

Once you've installed the drivers, the adapter will automatically detect whether a wireless connection is available. When I first installed the device on a Windows XP system without a wireless access point, XP reported that the device wasn't functioning because it didn't detect a wireless connection of any kind.

When I reinstalled later with an active wireless access point—I tested the adapter with the SMC7004WFW Barricade Plus Wireless Cable/DSL Broadband Router—I received no error message, and the adapter automatically detected the wireless connection. But it didn't connect immediately. It took a minute or two for the adapter to make a connection after the successful installation. Once it did, however, I was up and running on the network—able to access shared folders and surf the Internet via the broadband connection.

Performance

The literature accompanying the SMC2664W states that it has a working range of up to 422 feet at 11 Mbps and 825 feet at 1 Mbps. These numbers are probably dependent on having a clear path to the wireless access point. I unscientifically tested the range by taking my laptop outside and walking down the street until I finally lost the signal. With the Barricade sitting on the desk in my basement, I was able to maintain a network connection from about 150 feet from the house. A number of obstacles stood between the adapter and the router, including the concrete wall that encloses the basement. I was impressed that I was able to get as far down the road as I did before I lost the connection. If the neighbors on either side of me had adapters, they could be surfing the Internet on my broadband connection right now.

When you move the cursor over the connection icon in the system tray, a tool tip tells you how good your connection is. When I initially connected just a few feet away from the router, the tip said the connection was "excellent." When I moved upstairs, the connection went from "excellent" to "very good." The connection was just "good" when I moved outside until I lost the connection altogether.

After I lost the connection, it took a minute or two for the adapter to restore it when I moved back into range. Once you get just out of range, you have to move back to a point where you're receiving a strong signal and wait a short time before the adapter can restore the connection.

In one of the tests I performed, I downloaded a 5-MB file from Download.com to see if I could get the same kind of speed that I enjoy while directly connected via a Cat 5 cable. Surprisingly, the download speed was 342 KBps.


While I was copying a large file from the desktop PC I was connecting to, the screen-saver on the desktop system kicked in and the laptop promptly locked up, forcing a hard reboot. After the laptop came back up, I was able to retrieve the file without a problem. It took four minutes to copy the 100-MB file from the desktop to the laptop.

I found that several hours later when I turned on the laptop to reconnect, the adapter didn't detect that the wireless connection was available. I had to manually reconnect to the network. In Windows XP, this meant going into Wireless Network Connection Settings and selecting the name of the wireless connection. This restored the connection immediately, though.

Easy to use and quick to connect

The SMC2664W is a simple product that does what it's designed to do. For home or small

office users, it's a no-fuss way to get connected quickly to a wireless network. Even novice computer users can install it and get up and running in no time. The adapter is available online for as little as \$60, but retailers are charging \$75 and up for it.

The only obstacle I encountered was the floppy disk issue. I might be in the minority when it comes to users without floppy drives, but I think that will change. So it would be nice to see SMC offer their drivers on a CD. 

SMC's wireless broadband router offers performance tempered with caveats

Jan 27, 2003

By Ray Geroski

A variety of networking products aimed at small offices/home offices (SOHOs) have flooded the market, with huge leaps having been made in wireless networking. The intent is to make it easier for SOHO users to set up networks and share data, and, for the most part, the vendors have achieved that goal.

Among these offerings is SMC's 7004WFW Barricade Plus wireless broadband router. Aimed primarily at SOHO users, this 10/100 Mbps three-port router also acts as an 11-Mbps wireless access point, allowing clients with wireless adapters to share a broadband Internet connection. The 7004WFW is easy to set up and performs well, but some caveats make it a less-than-perfect product. On the whole, the 7004WFW is a good option, but you can find comparable products—even similar SMC offerings—that cost less.

Product details

Barricade Plus (**Figure A**) is packaged with the router unit, two antennae, power adapter, and setup CD. You must attach the antennae to the router by screwing them in place on the port-side of the box. The manual included with the device is merely a quick start guide, but you can download a more detailed document in PDF format from the SMC Web site. You really won't need either at first, however, because the setup CD entirely automates the initial setup process. It's only when you get into advanced settings—filtering and security measures—that you might need to consult the manual.

SMC lists the following specifications and features for Barricade Plus:

- ▶ IEEE 802.11b compliant
- ▶ Wireless operation at 11, 5.5, 2, or 1 Mbps
- ▶ Range of up to 304.8m (1,000 ft.)
- ▶ Frequency: (U.S./Canada/Europe) 2.400-2.4835 GHz; Japan: 2.471-2.497 GHz

- ▶ Internet access: 10/100 Mbps WAN port connection to xDSL/cable modem
- ▶ Home networking: Three 10/100 Mbps Ethernet switch ports with MDI/MDI-X autonegotiation
- ▶ Configurable parental control by limiting access to Web sites with URL and keyword blocking
- ▶ Stateful packet inspection (SPI) advanced firewall protection
- ▶ Client privileges, intrusion detection, and NAT
- ▶ Built-in VPN tunneling

Installation and performance

SOHO users will appreciate Barricade Plus's simple install process, which is accomplished via SMC's EZ 3-Click Installation Wizard (**Figure B**). True to its name, it takes but three clicks to complete the process. The program automatically detects the settings required for the broadband connection and assigns IP addresses to the internal network.

You can have your wireless network up and running in a matter of minutes. SMC claims a range of 1,000 feet for the wireless signal, but you'll find that walls greatly diminish this. Small offices will still have the freedom of networking from just about anywhere in the building, and home users will find they can obtain and maintain a connection without a problem—as long as they don't live in a mansion. In "Untether your network with SMC's wireless adapter" (page 176), we provide information on range tests with the SMC2664W wireless adapter. Results with different adapters will likely vary, but this should give you a good idea what you'll be able to do with Barricade Plus and whatever wireless adapters you're using. This article also includes details on wireless data transfer tests. As you'll find from other tests of the device, the SMC7004WFW performs well, holding its own in comparisons with similar devices.

One benefit of using Barricade Plus is that you don't have to pull any network cable to allow users to be able to communicate, share files, and access the Internet. It breaks down barriers that might otherwise prevent some

Figure A



SMC Barricade Plus wireless broadband router

Figure B



EZ 3-Click Installation Wizard

users from being able to freely access network resources; you can take your laptops to the conference room for meetings without having to deal with an octopus of cables. The range of Barricade Plus is good enough to make it an effective tool for such purposes. Obstacles tend to degrade the signal, but you'll still find that within the effective range you can maintain a reliable connection.

VPN

The feature that small businesses may appreciate most is that Barricade Plus can also serve as a VPN router. It represents a relatively inexpensive VPN solution, though many

Figure C

PPTP Tunnel Setting

This page contains the detail PPTP tunnel settings. You may specify the Idle Time Out which defines the idle time out period when PPP session will be terminated after the period of time without traffic going through. You may also configure the tunnel to behave as either client or server. No more than one PPTP client is allowed to be enabled at any time. For a client tunnel, both host mode and router mode (LAN-to-LAN) are supported. The tunnel can also be configured to automatically reconnect to the server when there is Internet traffic generated.

User name:	RayG			
password:	●●●●●●●●			
Idle Time Out:	15	(Min)		
IP:	192	168	4	0
Subnet Mask:	255	255	252	0
Gateway IP:	12	220	72	1
Client Setting:	<input checked="" type="checkbox"/> pptp client <input checked="" type="checkbox"/> host <input type="checkbox"/> auto reconnect			

OK Cancel

PPTP account details

Figure D

IPsec

IPsec allows users to define three secure IPsec tunnels with remote end points. This page includes the setting for both ends of each secure tunnel and the inbound/outbound Security Association (SA).

Enable IPsec: ☒ Yes ☐ No

Tunnel: Tunnel 1 ▼

Tunnel Enable : ☒ Yes ☐ No

Inbound SA (same as Outbound SA)

SPI: 256

Local IP Address: 192 168 2 0

Subnet Mask: 255 255 255 0

Remote IP Address: 192 168 3 0

Subnet Mask: 255 255 255 0

Security Gateway: 192 168 1 100

Hash Algorithm: ☐ None ☒ MD5(32 character) ☐ SHA1(40 character)

Key: 123456789abcdef02468ace01357

☐ None ☒ 3DES_CBC(48 character) ☐ DES_CBC(16 character)

Encrypt Algorithm: 0123456789abcdef02468ace1357

HELP APPLY CANCEL

IPSec tunnel setup

comparable products offer the same feature at a lower cost. The manual for setting up the VPN is not included with the package but is available for download at SMC's Web site (http://www.smc.com/drivers_downloads/library/7004xFW_VPN_QIG.pdf).

Setting up the VPN is fairly straightforward. **Figure C** shows the options for configuring the PPTP settings. You can create up to 20 PPTP VPN user accounts to manage remote access to the network.

You can also configure up to three IPSec tunnels. **Figure D** shows the interface for setting up an IPSec tunnel. Note the available encryption options for securing access.

After the user accounts are set up and the IPs configured, all users can set up the VPN connections on their remote systems and log in with their usernames and passwords.

Via the browser interface, you can quickly and easily give telecommuters remote access to your network. Given what typical VPN routers cost, this makes Barricade Plus an attractive option. But unfortunately, it might not make for the most reliable choice.

Caveats


One annoyance I encountered with Barricade Plus is that it would frequently lose the connection to the Internet. It operated flawlessly for about a month, and then I suddenly began to experience a rash of frequent disconnections. It was easy enough to restore the connection by accessing the router settings via the browser interface and releasing and renewing the configuration. Over time, though, I found myself having to do this more and more often. On one occasion, even after renewing the settings, the router was unable to reconnect. I had to completely disconnect it and turn it off to clear out the settings and then start over with the installation wizard.

Judging from other reports I've read scattered about the Internet, this is not an isolated experience. When I contacted SMC tech support about the issue, they sent me a beta firmware upgrade. But the upgrade failed to install correctly. The update utility could not detect the new file and, thus, could not install it.

When the SMC7004WFW works, it works well and offers many features useful to small businesses, including VPN support. It's those times when it doesn't work that make me

reluctant to recommend it wholeheartedly. The frequent downtime would certainly blunt its usefulness as a VPN solution, especially if the disconnects occurred during nonworking hours when no one could troubleshoot the problem. In small businesses without dedicated IT staffs, this could be a real hassle.

The price tag and availability of the device are also factors that may give you pause. I couldn't find many online vendors that currently offer the product, and the ones that do are selling it for over \$200. That's nearly double what many comparable products are currently going for. Even SMC's own product line includes similar products that cost less. If you can find the SMC7004WFW for around \$100, it would be a decent bargain, but there's no sense paying over \$100 for it when you can get better products for less.

The wireless broadband router market is getting pretty crowded these days, so the SMC7004WFW will have a tough time competing at its current price tag. And given the Internet connection issues, it may not be the best option for small businesses that need a reliable VPN. 

Vivato's WLAN switches extend Wi-Fi range

Jan 9, 2003

By Ray Geroski

Wireless networking is evolving at a rapid pace, transforming into an increasingly viable solution for enterprise networks. Performance boosts and security enhancements are among the forces pushing wireless deeper into the networking mainstream. Vivato is taking that evolution a step further with the introduction of indoor and outdoor Wi-Fi switches. Vivato expects to launch the products in early 2003. They'll be the first wireless switching products to hit the market and could have a significant impact on how organizations deploy wireless networks. A single switch installed indoors will offer network coverage for an entire floor, while an outdoor switch can connect buildings.

Service providers will be able to take advantage of the several-mile range of the outdoor switch, and the indoor switches' shorter range of around two miles will give organizations much more flexibility in deploying wireless networks by replacing multiple access points.

The technology Vivato is introducing could represent a significant step forward for wireless networking. Any organization that has implemented WLANs or is planning a WLAN implementation this year should take note.

Company and product background

Vivato is a two-year-old startup headquartered in San Francisco. Unlike other companies in the wireless market, Vivato is focusing on infrastructure products rather than client devices. To achieve wireless switching, Vivato combined existing smart antenna technology with existing wireless technology.

"People thought that combination was impossible," said Vivato Vice President of Marketing Phil Belanger, "but we figured out how to make it work, and we were able to accomplish it because we had multiple disciplines on our team."

An important part of the knowledge necessary to combine the two technologies came

from the founders' backgrounds in the cellular phone industry with Agilent. Belanger said that their experience building test equipment for the cellular phone industry and working with sophisticated RF equipment was instrumental in their work to marry smart antennas with wireless technology.

Belanger said that Vivato's switches represent a big change in wireless networking products largely because, up until now, most of the devices introduced have been client adapters and access points. Belanger sees wireless switching as a sign that the market is maturing and believes that the introduction of these products represents a new architecture. He said that this evolution of wireless networking is analogous to that of Ethernet networking.

"When Ethernet switching was introduced, it helped the explosion of Ethernet because it really scaled up the capacity so it could work in large installations."

Another parallel is in the constant upgrading of the speed of Ethernet. Belanger said the same kind of thing is happening in wireless networking. The market is beginning to explode because of the rapid improvements being made.

Belanger feels that Vivato's switches offer a more robust way to deploy wireless networks with a lower TCO because there will be fewer pieces of hardware to install. For example, organizations currently have to deploy a number of access points to provide network coverage on a single floor. But with the Vivato switches, they'll deploy just one device on a building floor to achieve the same connectivity.

The switches will operate much like standard gigabit Ethernet switches, and the devices include support for VPNs, VLANs, and 802.1X security. Because the switches will operate in a familiar manner and because they will replace many devices that would otherwise have to be installed, Belanger said that the Vivato wireless solution will be easier to manage.

The range of the outdoor switch Vivato has developed also makes it a solution well tailored for broadband delivery.

“We have phenomenal range. We’ve taken the range of Wi-Fi from something that operates at tens of meters to something that operates at kilometer distances.”

Belanger said from the service-provider perspective, this solution presents a potential alternative for delivering broadband to the home more cost effectively. Instead of a DSL or cable modem/router that costs \$100 or more to act as an Internet gateway, Belanger said the gateway in users’ homes could be a Wi-Fi card (or a newer notebook with Wi-Fi built in). The cost of the client equipment would be lower, and it would be much easier to implement.

Performance, reliability, and security

Although wireless networking may not yet be an actual substitute for wired networking, it does offer some benefits that make it attractive in many cases. Where freedom of mobility is important, wireless networking has obvious advantages. And because it does not require pulling cables or making accommodations for wiring, it can potentially lower infrastructure costs, especially in locations with certain limitations. At the very least, wireless networking is easier to install because of the absence of cabling.

Although wireless networking is not as robust in terms of performance as wired networking, Belanger pointed out that the evolving standards for wireless networking are at least making it as reliable. However, he noted that the measures taken to make wireless reliable have come with some overhead costs in terms of performance.

“To get to that same level of robustness, there’s overhead. The cost of building in the reliability is that it’s not as efficient as Ethernet. Out of the 11 MB provided for in 802.11b, you might get 6 1/2 MB of actual throughput.”

The other question about wireless networking—especially with Vivato’s switches, which have an increased range—is security. Belanger

said that along with the security built in to the boxes, including support for the Wi-Fi Alliance-backed Wireless Protected Access (WPA) protocol, VPN, low-level encryption, and 802.1X authentication and key distribution, the Vivato switches offer the security benefit of using focused beams rather than a broadcast signal.

“The switches aren’t putting out any more real energy than a conventional access point. Most of the range improvement comes from antenna gain that allows us to set up very narrow beams of Wi-Fi.”

The narrowing of the beam, Belanger said, makes it possible for the switches to achieve the greater range. The security benefit is that the signal isn’t being broadcast in all directions to be easily intercepted. Beams are focused where they need to go (and when they need to go) rather than broadcasting all the time.

“We can point the beams precisely at the intended clients, and we can move them around on a packet-by-packet basis.... Unless you’re close in, you’re not going to hear much traffic. You’ll pick up traffic only when the beam is aimed in your direction,” Belanger said.

Through a built-in scanning function that operates all the time, the switches can locate the clients and direct the beam. Once the function finds a client, it records information about the location so it knows where to send transmissions.

“We use that function to collect information about the nodes we want to participate on the network, but we can use that same function to detect rogue access points.”

The software that’s shipping with the device will allow users to take advantage of the security capabilities of the switch. Another security benefit is that the onboard hardware accelerator supports 802.11i, which changes the encryption method and will require hardware upgrades for many wireless products. Vivato’s switches will be firmware upgradeable to support the new standard.

Bottom line

Because Vivato’s switches introduce new wireless functionality and incorporate a number of useful features, they could represent a

big step forward for wireless networking in the enterprise. The switches may eliminate the need for many wireless devices currently in use and extend the effective range of wireless connectivity, and they offer built-in secu-

rity features. Organizations should pay close attention to wireless switching when Vivato's products arrive. ~

Notes

Notes

Notes

Builder.com | CNET.com | TechRepublic.com | ZDNet.com

e-mail: customerservice@techrepublic.com

Phone: 845-566-1866 • 800-217-4339

Product code: B059

