W LAN TROUBLESHOOTING

DAVID D. COLEMAN, **CWNE #4**, AEROHIVE... A Wiley Brand

WLAN Troubleshooting

Excerpt from Certified Wireless Network Administrator Official Study Guide Exam CWNA-107



David D. Coleman, CWNE #4



Development Editor: Kim Wimpsett Technical Editor: Ben Wilson Copy Editor: Judy Flynn Editorial Manager: Pete Gaughan Production Manager: Kathleen Wisor Executive Editor: Jim Minatel Proofreader: Nancy Carrasco Project Coordinator, Cover: Brent Savage Cover Design: Brendan Bond

Copyright © 2018 by John Wiley & Sons, Inc., Indianapolis, Indiana

Published simultaneously in Canada

ISBN: 978-1-119-37310-0 ISBN: 978-1-119-37312-4 (ebk.)

Manufactured in the United States of America

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 646-8600. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at http://www.wiley.com/go/permissions.

Limit of Liability/Disclaimer of Warranty: The publisher and the author make no representations or warranties with respect to the accuracy or completeness of the contents of this work and specifically disclaim all warranties, including without limitation warranties of fitness for a particular purpose. No warranty may be created or extended by sales or promotional materials. The advice and strategies contained herein may not be suitable for every situation. This work is sold with the understanding that the publisher is not engaged in rendering legal, accounting, or other professional services. If professional assistance is required, the services of a competent professional person should be sought. Neither the publisher nor the author shall be liable for damages arising herefrom. The fact that an organization or Web site is referred to in this work as a citation and/or a potential source of further information does not mean that the author or the publisher endorses the information the organization or Web site may provide or recommendations it may make. Further, readers should be aware that Internet Web sites listed in this work may have changed or disappeared between when this work was written and when it is read.

For general information on our other products and services visit www.wiley.com/go/custompub.

TRADEMARKS: Wiley, the Wiley logo, and the Sybex logo are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates, in the United States and other countries, and may not be used without written permission. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc. is not associated with any product or vendor mentioned in this book.

10987654321

About the Author

David D. Coleman is the Senior Product Evangelist for Aerohive Networks, www.aerohive.com. David collaborates with the Aerohive Product Training team and travels the world for WLAN training sessions and speaking events. He has instructed IT professionals from around the globe in wireless networking administration, wireless security, and wireless frame analysis. David has written multiple books, blogs, and white papers about wireless networking. Prior to working at Aerohive, he specialized in corporate and government Wi-Fi training. In the past, he provided WLAN training for numerous private corporations, the US military, and other federal and state government agencies. When he is not traveling, David resides in Atlanta, Georgia. David is CWNE #4, and he can be reached via email at mistermultipath@gmail.com. Please follow David on Twitter: @mistermultipath.

Acknowledgments

I would like to thank the entire Aerohive Networks Product Training department: Gregor Vucajnk, Marko Tisler, Chris Harkins and Henry Vu. Additionally, I need to thank Aerohive employees John Merrill, Abby Strong and Alan Amrod for their leadership this past year. Also special thanks to Wendy Kastner for her mockups. A very special shout-out must go to Aerohive's Brendan Bond for the awesome cover design he created for this booklet. There are many other past and present Aerohive employees I would like to thank, but there simply is not enough room. So thank you to all of my Aerohive co-workers. It has been an honor working with you to help building something special.

I also need to thank my co-author and good friend of many years, David Westcott. Although his name is not on this booklet, he contributed much of the content. I was going to ask him fully co-author this booklet but I could not locate him. #whereiswestcott

I need to send thanks to Joel Crane of both MetaGeek and Ekahau fame for his contributions and the spectrum analyzer screenshots. Several other WLAN rock stars are mentioned in the copy of this book: Mike Albano, Eddie Forero, James Garringer, Adrian Granados, Jerome Henry and François Vergès.

I would like to thank my development editor, Kim Wimpsett, who has been a pleasure to work with as always. Much love goes out to my technical editor, Ben Wilson. I also need to send special thanks to copy editor Judy Flynn, production manager Kathleen Wisor and proofreader Nancy Carrasco. Many thanks also to executive editor Jim Minatel and editorial manager, Pete Gaughan. And final thanks to Karen Hattan for her patience.



WLAN Troubleshooting

IN THIS BOOKLET, YOU WILL LEARN ABOUT THE FOLLOWING:

✓ Five Tenets of WLAN Troubleshooting

- Troubleshooting best practices
- Troubleshoot the OSI model
- Most Wi-Fi problems are client issues
- Proper WLAN design reduces problems
- WLAN always gets the blame

✓ Layer 1 troubleshooting

- WLAN design
- Transmit power
- RF interference
- Drivers
- PoE
- Firmware Bugs

✓ Layer 2 troubleshooting

- Layer 2 retransmissions
- RF interference
- Low SNR
- Adjacent channel interference
- Hidden node



- Mismatched power
- Multipath

✓ Security Troubleshooting

- PSK troubleshooting
- 802.1X/EAP troubleshooting
- VPN troubleshooting
- ✓ Roaming troubleshooting
- ✓ Channel utilization
- ✓ Layers 3-7 troubleshooting
- ✓ WLAN troubleshooting tools
 - WLAN discovery applications
 - Spectrum analyzers
 - Protocol analyzers
 - Throughput test tools
 - Standard IP network test utilities
 - Secure shell
- ✓ Aerohive troubleshooting tools



As with any type of communications network, problems with WLAN networks arise that

might require attention from an administrator. Client connectivity issues often arise that might be the result of improper implementation of WLAN security. In this booklet, you will learn about troubleshooting best practices and how to focus on layer 1 and layer 2 when investigating Wi-Fi problems. Being able to investigate roaming problems as well as monitor channel utilization is key to maintaining good WLAN health. You will also learn WLAN troubleshooting strategies from a security perspective. Although this booklet is not meant to be a step-by-step diagnostic guide, you will learn about many common WLAN problems and suggested resolutions. We will also discuss many of the freeware and commercial WLAN troubleshooting tools that are available. Finally, we will review some of the most useful Aerohive diagnostic tools and CLI commands.

Five Tenets of WLAN Troubleshooting

Before we discuss specific WLAN troubleshooting strategies, you should understand five basic tenets for troubleshooting any type of WLAN problem:

- Implement troubleshooting best practices.
- Troubleshoot the OSI model.
- Most problems are client side.
- Proper WLAN design/planning is important.
- The WLAN will always get the blame.

We will now review these WLAN troubleshooting doctrines in greater detail.

Troubleshooting Best Practices

The fundamentals of troubleshooting best practices are to ask questions and collect information. When troubleshooting any type of computer network,

4 WLAN Troubleshooting

you must ask the correct questions to collect information that is relevant to the problem. It is easy to get sidetracked when troubleshooting, so asking the proper questions will help an IT administrator focus on the pertinent data with a goal of isolating the root cause of the problem. For example, WLAN security problems often result in WLAN client connectivity issues; asking the appropriate questions will point you in the right direction toward solving the problem. The following basic questions are among those that need to be asked:

When is the problem happening?

At what time did the problem occur? Did this problem happen during a very specific time period? This information can be easily determined by looking at the log files of access points (APs), and applicable servers such as RADIUS. Best practices mandate that all *Network Time Protocol (NTP)* and time zone settings be correctly configured on all network hardware.

Where is the problem happening?

Is the problem widespread or does it only exist in one physical area? Is the problem occurring on a single floor or in the entire building? Does the problem affect just one access point or a group of access points? Determining the location of the problem will help you gather better information toward solving the problem.

Does the problem affect one client or numerous clients?

If the problem is only affecting a single client, you may have a simple driver issue or an incorrectly configured supplicant. If the issue is affecting numerous clients, then the problem is obviously of greater concern. Most connectivity problems are client side, whether they are detrimental to a single client or multiple clients.

Does the problem reoccur or did it just happen once?

Troubleshooting a problem that only happens one time or only a few times can be difficult. Collecting data is much easier with recurring problems. You may have to enable debug commands on APs to hopefully capture the problem again in a log file.

Did you make any changes recently?

This is a question that the support personnel of WLAN vendors always ask their customers. And the answer is almost always no despite the fact that changes to the network indeed take place. Best practices dictate that any network configuration changes be planned and scheduled. WLAN infrastructure security audit logs will always leave a paper trail of which administrator made which changes at any specific time.

Once you have asked numerous questions, you can begin the process of solving the problem. Troubleshooting best practices include the following:

1. Identify the issue.

Because the WLAN always seems to get the blame, it is even more important to correctly identify the problem. Determine that a problem actually exists. Asking questions and collecting information will help you identify the true issue.

2. Re-create the problem.

Having the ability to duplicate the problem either onsite or in a remote lab gives you the ability to collect more information to diagnose the problem. If you cannot re-create a problem, you may need to ask more questions.

3. Locate and isolate the cause.

The whole point of asking the pointed questions and gathering data is so that you can isolate the root cause of the problem. Troubleshooting up the OSI model will also help you identify the culprit. Identify if the problem is at the access layer, distribution layer, or core of your network design.

4. Solve the problem.

Formulate and implement a plan to solve the problem. This may require network changes, firmware updates, and so forth.

5. Test to verify the problem is solved.

Always be sure to test in different areas during different times and with multiple devices. Extensive testing will ensure that the problem is indeed resolved.

6. Document the problem and the solution.

Troubleshooting best practices dictate that you document all problems, diagnostics, and resolutions. A reference help desk database will assist you in solving problems in a timely fashion should any problem reoccur.

7. Provide feedback.

As a professional courtesy, always be sure to follow up with the individual(s) who first alerted you to the problem.

Troubleshoot the OSI Model

The diagnostic approach that is used to troubleshoot wired 802.3 networks should also be applied when troubleshooting a wireless local area network (WLAN). A bottoms-up approach to analyzing the OSI reference model

layers also applies to wireless networking. Remember that 802.11 technology is similar to 802.3 in that it operates at the first two layers of the OSI model. For that reason, a WLAN administrator should always try to first determine whether problems exist at layer 1 and layer 2. If the first two layers of the OSI model have been eliminated as the cause of the problem, the problem is not a Wi-Fi problem and the higher layers of the OSI model should be investigated.

As with most networking technologies, most problems usually exist at the Physical layer. Simple layer 1 problems, such as nonpowered access points or client radio driver problems, are often the root cause of connectivity or performance issues. Disruption of RF signal propagation and RF interference will affect both the performance and coverage of your WLAN. Inadequate WLAN coverage, capacity, and performance are often layer 1 problems that are a result of poor WLAN design. Client driver issues and misconfigured supplicants are also common layer 1 problems.

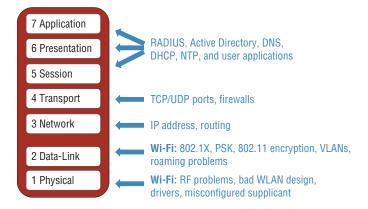
After eliminating layer 1 as the source of the problem, a WLAN administrator should try to determine whether the problem exists at the Data-Link layer. Basic 802.11 communications such as discovery, authentication, association, and roaming all occur at the MAC sublayer of layer 2. As shown in Figure 1, WLAN security mechanisms also operate at layer 2. Modern-day 802.11 radios use CCMP encryption that provides data privacy for layers 3-7. The chosen encryption method must match on both the AP and client radios. For example, if an AP has disabled backward compatibility for TKIP encryption, a legacy client that only supports TKIP will not be able to connect. Remember that only CCMP encryption can be used for 802.11n (HT) and 802.11ac (VHT) data rates. An access point might be configured to transmit an SSID that supports both TKIP and CCMP encryption. In this situation, a common support call may be that the legacy TKIP clients seem slow because of the lack of support for higher data rates. The simple solution is to replace the legacy clients with modern-day clients that support CCMP.

There is a symbiotic relationship between the creation of dynamic encryption keys and authentication. A pairwise master key (PMK) is used to seed the 4-Way Handshake that generates the unique dynamic encryption keys employed by any two 802.11 radios. The PMK is generated as a byproduct of either PSK or 802.1X/EAP authentication. Therefore, if authentication fails, no encryption keys are generated. We will discuss troubleshooting both 802.11 authentication methods later in this booklet.

As stated earlier, if the first two layers of the OSI model have been eliminated, the problem is not a Wi-Fi problem, and therefore the problem exists within layers 3–7. It is likely the problem is either a TCP/IP

networking issue or an application issue. As shown in Figure 1, TCP/IP problems should be investigated at layers 3–4, whereas most application issues exist between layers 5 and 7.

FIGURE 1 OSI model



Most Wi-Fi Problems Are Client Issues

As previously mentioned, whenever you troubleshoot a WLAN, you should start at the Physical layer. Additionally, 70 percent of the time the problem will reside on the WLAN client. If there are any client connectively problems, WLAN Troubleshooting 101 dictates that you disable and re-enable the WLAN network adapter. The driver for the WLAN network interface card (NIC) is the interface between the 802.11 radio and the operating system (OS) of the client device. For whatever reason, the WLAN driver and the OS of the device may not be communicating properly. A simple disable/re-enable of the WLAN NIC will reset the driver. Always eliminate this potential problem before investigating anything else. Additionally, first-generation radio drivers and firmware are notorious for possible bugs. Always make sure the WLAN client population has the latest available drivers installed. Another change that is quick and easy to make is to reconfigure the client configuration profile. Most client supplicants allow the user to define a WLAN configuration profile or connection parameters. Sometimes troubleshooting a problem is as easy as deleting the old profile and configuring a new profile.

As mentioned earlier, client-side security issues usually evolve around improperly configured supplicant settings. This could be something as simple as a mistyped WPA2-Personal passphrase or as complex as 802.1X/EAP digital certificate problems.



Is There a Master Database of Wi-Fi Client Capabilities?

The short answer is that there is not any official database of 802.11 client devices and their capabilities. There are, however, a few resources, including the Wi-Fi Alliance, which maintains a Wi-Fi CERTIFIED Product Finder database at www.wi-fi.org/product-finder. Although most WLAN infrastructure vendors submit their access points for certification, it should be understood that many manufacturers of WLAN client devices do not go through the certification process. As shown in Figure 2, Wi-Fi expert Mike Albano (CWNE #150) maintains a free public listing of WLAN client capabilities at clients.mikealbano.com. Mike has put together a good database of many of the modern-day popular WLAN client devices. You can also download 802.11 frame captures of the client devices as well as submit WLAN client information. Often, a laptop or mobile device manufacturer will list the radio model in the specification sheet for the laptop or mobile device. However, some manufacturers may not list detailed radio specifications and capabilities. Another method of identifying the Wi-Fi radio in your device is from the FCC ID. In the United States, all Wi-Fi radios must be certified by the Federal Communications Commission (FCC) government agency. The FCC maintains a searchable equipment authorization database at www.fcc.gov/fccid. You can enter the FCC ID of your device into the database search engine and find documentation and pictures submitted by the manufacturer to the FCC. The FCC database is very useful in helping identity Wi-Fi radio models and specifications if the information is not available on the manufacturer's website.

FIGURE 2 WLAN client database



Proper WLAN Design Reduces Problems

A huge percentage of WLAN support phone calls are a symptom of a lack of WLAN design. Proper capacity and coverage planning, spectrum analysis, and a validation site survey will eliminate the majority of WLAN support tickets in regard to performance. Proper WLAN design will minimize problems such as co-channel interference (CCI) in advance. Additionally, many WLAN security holes can be eliminated in advance with proper WLAN security planning. If 802.1X/EAP is deployed, one of the biggest challenges is how to provision the root CA certificates for mobile devices such as smartphones and tablets. A well-thought-out security strategy for employee WLAN devices, BYOD devices, and guest WLAN access is essential. Proper WLAN planning and design in advance will reduce time spent troubleshooting WLAN problems at a later juncture.

WLAN Always Gets the Blame

Despite all your best WLAN troubleshooting practices and best efforts, you should resign yourself to the fact the WLAN will always get the blame. Experienced WLAN administrators know the WLAN will be blamed for problems that have nothing to do with the WLAN. This is another reason that troubleshooting up the OSI stack is important. If the problem is not a layer 1 or layer 2 issue, then Wi-Fi is not the culprit. However, put yourself in the shoes of the end user who is connected to the WLAN. 802.11 technology exists at the access layer. The whole point of an AP is to provide a wireless portal to a preexisting network infrastructure. Your employees and guests who connect to the WLAN expect seamless wireless mobility, and they have no concept of problems that exist at layers 3–7. A WLAN end user is not aware that the DHCP server is out of leases. A WLAN end user is not aware the Internet service provider (ISP) is experiencing difficulty and the WAN link is down. The WLAN end user just knows that they cannot access www.facebook.com through the WLAN and therefore they point the finger at the Wi-Fi network.

Layer 1 Troubleshooting

As discussed in the beginning of this booklet, most networking problems usually exist at the Physical layer. In the following sections, we will delve deeper into layer 1 problems often caused by bad WLAN design or RF interference. We will also discuss layer 1 issues such as radio driver issues, firmware bugs, and issues related to Power over Ethernet (PoE).

WLAN Design

As previously mentioned, the bulk of WLAN problems can be avoided with good WLAN design prior to deployment. Probably the two most common layer 1 problems that arise due to poor design are coverage holes and cochannel interference. WLAN coverage holes are usually a result of a lack of a site survey validation. Confirmation of –70 dBm coverage for high data rate connectivity and –65 dBm coverage for voice-quality WLANs is imperative. Always remember that the AP radios' receive sensitivity is usually much stronger than client receive sensitivity. Measurements to validate received signal strength should be from the perspective of the client devices. Because of wide variances in client radio RSSI, the least sensitive client device is often used for validation of proper signal strength. A lower than desired received signal will result in radios shifting to lower data rates, which consume more airtime and negatively impact performance. Coverage dead zones often arise post-deployment when furniture and even walls are moved.

Bad WLAN coverage often is a result of improper placement of APs as well as improper antenna orientation. Always check the technical specifications of your WLAN vendor; however, most indoor access points with low-gain omnidirectional antennas should be mounted on the ceiling no higher than 3 meters from the floor. When using external omnidirectional antennas, it is important that they be vertically positioned. A common mistake is to position them horizontally. You would be shocked at how many APs are improperly installed, often with the APs mounted in the plenum with the antennas facing upward. You can view many pictures of improper AP placement at https://badfi.com/bad-fi, which is a fun blog written by Eddie Forero, CWNE #160.

Co-channel interference (CCI) is the top cause of needless airtime consumption that can be minimized with proper WLAN design best practices. Carrier Sense with Multiple Access Collision Avoidance (CSMA/CA) dictates half-duplex communications, and only one radio can transmit on the same channel at any given time. An 802.11 radio will defer transmissions if it hears the PHY preamble transmissions of any other 802.11 radio at an SNR of just 4 DB or greater. Unnecessary medium contention overhead that occurs when too many APs and clients hear each other on same channel is called *co-channel interference (CCI)*. In reality, the 802.11 radios are operating exactly as defined by the CSMA/CA mechanisms, and this behavior should really be called *co-channel cooperation*. However, the unnecessary medium contention overhead caused by co-channel interference is usually a result of improper channel reuse design. While it is almost impossible to prevent CCI in the 2.4 GHz band, the airtime consumption that is a result of CCI can be minimized and possibly avoided with good 5 GHz WLAN design

best practices. Making good use of proper 5 GHz channel reuse patterns and enabling the *dynamic frequency selection (DFS)* channels will reduce CCI. Lower transmit power will also reduce CCI.

Transmit Power

Another common layer 1 problem is the mistake often made when deploying access points, having the APs transmit at full power. Although most indoor APs may have full transmit power settings as high as 100 mW, they should rarely be deployed at full power. Effectively, this extends the effective range of the access point; however, designing WLANs strictly for coverage is an outdated concept. WLAN capacity and reducing airtime consumption are higher priorities. AP's at maximum transmit power will result in oversized coverage but not meet your capacity needs. Access points at full power will also increase the odds of co-channel interference due to bleed-over transmissions. Here is a quick summary of all the problems caused by APs at maximum power:

- Capacity needs not met
- Increase in CCI and airtime consumption due to unnecessary medium contention overhead
- Increase in hidden node issues
- Increase in sticky clients and roaming problems

For all of these reasons, typical indoor WLAN deployments are designed with the APs set at about one-fourth to one-third transmit power. Higher user density environments may require that the AP transmit power be set at the lowest setting of 1 mW.

RF Interference

RF interference from non-802.11 transmitters is by far the most common external cause of WLAN problems that exist at layer 1. The *energy detect* (ED) thresholds for non-802.11 transmissions are much higher than the *signal detect* (SD) threshold for detecting 802.11 radio transmissions. However, various types of RF interference can still greatly affect the performance of an 802.11 WLAN. Interfering devices may exceed the energy detect threshold and prevent an 802.11 radio from transmitting, thereby causing a denial of service. If another RF source is transmitting with strong amplitude, 802.11 radios can sense the RF energy during the *clear channel assessment* (CCA) and defer transmission entirely. The other typical result of RF interference is that 802.11 frame transmissions become corrupted. If frames are corrupted due to RF interference, excessive retransmissions will occur, and therefore

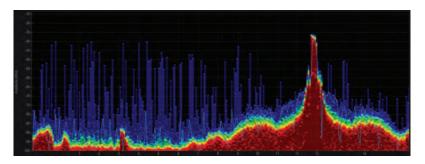
throughput will be reduced significantly. There are several different types of RF interference, as described in the following sections.

Narrowband Interference

A narrowband RF signal occupies a smaller and finite frequency space and will not cause a *denial of service* (*DoS*) for an entire band, such as the 2.4 GHz ISM band. A narrowband signal is usually very high amplitude and will absolutely disrupt communications in the frequency space in which it is being transmitted. Narrowband signals can disrupt one or several 802.11 channels.

Narrowband RF interference can also result in corrupted frames and layer 2 retransmissions. The only way to eliminate narrowband interference is to locate the source of the interfering device with a spectrum analyzer and remove the interfering device. To work around interference, use a spectrum analyzer to determine the affected channels and then design the channel reuse plan around the interfering narrowband signal. Figure 3 shows a spectrum analyzer capture of a narrowband signal close to channel 11 in the 2.4 GHz ISM band.

FIGURE 3 Narrowband RF interference



Wideband Interference

A source of interference is typically considered wideband if the transmitting signal has the capability to disrupt the communications of an entire frequency band. Wideband jammers exist that can create a complete DoS for the 2.4 GHz ISM band. The only way to eliminate wideband interference is to locate the source of the interfering device with a spectrum analyzer and remove the interfering device. Figure 4 shows a spectrum analyzer capture of a wideband signal in the 2.4 GHz ISM band with average amplitude of –70 dBm, well above the defined energy detect thresholds of all 802.11 radios.

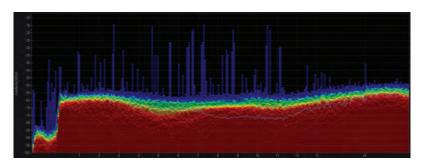


FIGURE 4 Wideband RF interference

All-Band Interference

The term *all-band interference* is typically associated with frequency hopping spread spectrum (FHSS) communications that usually disrupt the 802.11 communications at 2.4 GHz. FHSS constantly hops across an entire band, intermittingly transmitting on very small subcarriers of frequency space. A legacy 802.11 FHSS radio, for example, transmits on hops that are 1 MHz wide in the 2.4 GHz band. 802.11b radios transmit in a stationary 22 MHz of frequency space and 802.11 g/n radios transmit on fixed channels of 20 MHz of spectrum. While hopping and dwelling, an FHSS device will transmit in sections of the frequency space occupied by an 802.11b/g/n channel. Although an FHSS device will not typically cause a denial of service, the frame transmissions from the 802.11b/g/n devices can be corrupted from the all-band transmissions of a legacy 802.11 FHSS interfering radio.



Real World Scenario

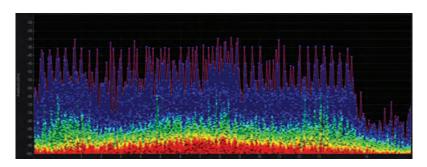
What Devices Cause RF Interference?

Numerous devices, including cordless phones, microwave ovens, and video cameras, can cause RF interference and degrade the performance of an 802.11 WLAN. The 2.4 GHz ISM band is extremely crowded, with many known interfering devices. Interfering devices also transmit in the 5 GHz U-NII bands, but the 2.4 GHz frequency space is much more crowded. The tool that is necessary to locate sources of RF interference is a spectrum analyzer. Luckily, most enterprise WLANs require the use of the 5 GHz band, which is less crowded and has much more frequency space. However, 5 GHz RF interferers also exist and spectrum analysis monitoring is needed.

Bluetooth (BT) is a short-distance RF technology used in WPANs. Bluetooth uses FHSS and hops across the 2.4 GHz ISM band at 1,600 hops per second. Older Bluetooth devices were known to cause severe allband interference. Newer Bluetooth devices utilize adaptive mechanisms to avoid interfering with 802.11 WLANs. Bluetooth adaptive frequency hopping is most effective at avoiding interference with a single AP transmitting on one 2.4 GHz channel. If multiple 2.4 GHz APs are transmitting on channels 1, 6, and 11 in the same physical area, it is impossible for the Bluetooth transmitters to avoid interfering with the WLAN. Digital Enhanced Cordless Telecommunications (DECT) cordless telephones also use frequency hopping transmissions. A now-defunct WLAN technology known as HomeRF also used FHSS; therefore, HomeRF devices can potentially cause all-band interference. Other frequency hopping devices that you may run across include various types of medical telemetry units. Although all the FHSS interferers mentioned so far transmit in the 2.4 GHz ISM band, 5 GHz frequency hopping transmitters that can cause interference also exist.

Frequency hopping transmitters do not usually result in as much data corruption as fixed-channel transmitters; however, the existence of a high number of frequency hopping transmitters in a finite space can result in a high amount of 802.11 data corruption and is especially devastating to VoWiFi communications. The only way to eliminate all-band interference is to locate the interfering device with a spectrum analyzer and remove the interfering device. Figure 5 shows a spectrum analyzer capture of a frequency hopping transmission in the 2.4 GHz ISM band. After locating the sources of interference, the best and simplest solution is to eliminate them entirely.

FIGURE 5 All-band RF interference



Drivers

As mentioned previously, first generation drivers for client devices often cause connectivity and roaming problems. Always check with the client device manufacturer to make sure you are using the most up-to-date drivers. Another thing to keep in mind is backward compatibility between newer access points and older client devices. Although the 802.11 amendments make provisions for backward compatibility, the opposite is often true in the real world. Legacy client drivers do not know how to handle the new fields in 802.11 information elements found in beacon and other management frames transmitted by an AP. When new technology is enabled on an AP, legacy clients often can no longer connect.

For example, roaming and connectivity problems may be a direct result of lack of support for 802.11k/v/r mechanisms on the client. Figure 6 depicts the two information elements that are seen in management frames sent by an AP with 802.11r enabled. The *mobility domain information element (MDIE)* and the *fast BSS transition information element (FTIE)* are fields of information necessary for APs and clients that support Voice Enterprise roaming capabilities. The drivers of legacy clients that do not support 802.11r may ignore these information fields and everything will be fine. But the legacy client drivers may also be disrupted by the 802.11r information elements, and client connectivity problems occur.

FIGURE 6 Fast BSS transition information element

```
Tag: Mobility Domain
  Tag Number: Mobility Domain (54)
  Tag length: 3
  Mobility Domain Identifier: 0x3b4d
  FT Capability and Policy: 0x00
  .... ...0 = Fast BSS Transition over DS: 0x0
  .... ..0. = Resource Request Protocol Capability: 0x0
Tag: Fast BSS Transition
  Tag Number: Fast BSS Transition (55)
  Tag length: 101
  MIC Control: 0x0000
  0000 0000 .... = Element Count: 0
  Subelement ID: PMK-R1 key holder identifier (R1KH-ID) (1)
  Length: 6
  PMK-R1 key holder identifier (R1KH-ID): 08ea4476b568
  Subelement ID: PMK-R0 key holder identifier (R0KH-ID) (3)
  Length: 9
  PMK-R0 key holder identifier (R0KH-ID): AH-76b540
```

Most businesses and corporations can eliminate many of the client connectivity and performance problems by simply upgrading company-owned client devices before updating the WLAN infrastructure. Sadly, the opposite is often more common, with companies spending many hundreds of thousands of dollars on technology upgrades with new access points while still deploying legacy clients.

PoE

WLAN vendors commonly receive support calls from customers complaining that all of a sudden access points randomly begin to reboot. In most cases, the root cause of random rebooting of APs is that the switch power budget has been eclipsed. Very often, if an AP cannot get the power that it needs, the AP will reboot and try again. The power budget of a switch or multiple switches should be monitored to make sure that all devices can maintain power. Active power budget information can usually be seen from the command line of a switch or the GUI interface or monitored by a centralized network management server (NMS).

Although proper power budgeting can prevent this problem during the design phase, remember that other devices, such as desktop VoIP phones, also use PoE. An extra PoE-powered device might have been plugged into a switch port and the power budget has been exceeded. Proper power budgeting and monitoring for access points and any other PoE-capable devices is paramount. Power budget problems will grow with the introduction of more 4x4:4 access points which will require more than the 15.4 Watts defined by 802.3af. As we move toward the next generation of 802.11ax access points, the use of 802.3at power will be a necessary requirement.

PoE provided by switches can also be your friend when trying to trouble-shoot an AP that for whatever reason may be unreachable from a remote location. For example, an AP may be inoperable due to some sort of processor overload and can no longer be monitored from an NMS or reachable via SSH. A simple forced reboot of the AP may restore communications. One of the oldest tricks used by network administrators is to *power cycle* the PSE port of the access switch that is providing power to the unresponsive AP. Enabling and disabling the power will force a reboot of the AP that may be locked up.

Firmware Bugs

As mentioned previously, older client firmware and drivers often causes connectivity issues with newer model APs. Conversely, updating access points with new firmware may also result in unexpected WLAN connectivity and more often performance problems. As with any type of networking device, an upgrade of the AP operating system is often needed when WLAN

vendors introduce new features and capabilities. Prior to the release of the newer AP code, WLAN vendors perform *regression testing*, which verifies that previously developed features and capabilities still operate and perform in the same manner. Despite the regression testing, new performance bugs may arise once the newer firmware is deployed in the field in enterprise environments.

A suggested deployment best practice would be to upgrade APs in a staging area for testing prior to wide scale deployment. Another strategy would be to update the APs in one building with active clients and see if any new problems arise. Once the firmware has been validated to be stable, a full upgrade of all the company APs can occur. Larger enterprise companies very often have clearly defined change management processes for making any type of network change, including firmware updates. Whenever possible, proper testing should also take place before the introduction of new client devices to the company WLAN.

When troubleshooting possible bugs, it will be necessary to engage with the support personnel of your WLAN vendor. You will most likely be requested to supply tech data logs and possibly packet captures. Many WLAN vendors may also offer an older version of firmware that has been thoroughly field-tested in enterprise environments and is often considered as the best bet in terms of quality assurance. Please understand that older versions of AP firmware may not offer you the new feature bells and whistles that you desire. Another advantage of updating APs to newer firmware is that previously discovered bugs may be fixed in the new release.

If a new bug is discovered once your APs are updated, the WLAN vendor's support team may recommend that you roll back your APs to the previous version of code until the bug can be addressed. When updating any access points, or other networking devices, always read the release notes to verify new features, fixed bugs, and known issues.

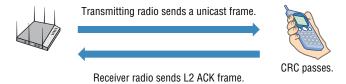
Layer 2 Troubleshooting

Wi-Fi radios communicate via 802.11 frame exchanges at the MAC sublayer of the Data-Link layer. Therefore, the next logical layer to troubleshoot in the OSI model is layer 2. The following sections will cover in great detail the many causes of layer 2 retransmissions and the significant adverse effects they cause. When troubleshooting a WLAN, you will always monitor layer 2 retransmission metrics.

Layer 2 Retransmissions

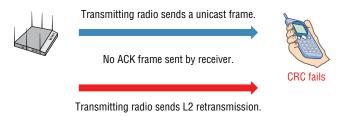
The mortal enemy of WLAN performance is layer 2 retransmissions that occur at the MAC sublayer. As shown in Figure 7, all unicast 802.11 frames must be acknowledged. In the trailer of each frame is a *cyclic redundancy check (CRC)*. The receiver 802.11 radio uses the CRC of the frame to confirm the data integrity of the payload of the incoming frame. If the CRC passes, the frame has not been corrupted during transit. The receiver 802.11 radio will then send an 802.11 acknowledgment (ACK) frame back to the original transmitter. Layer 2 ACK frames are used as a method of delivery verification.

FIGURE 7 Layer 2 ACK



If a collision occurs or any portion of a unicast frame is corrupted, CRC will fail, and the receiving 802.11 radio will not return an ACK frame to the transmitting 802.11 radio. As shown in Figure 8, if an ACK frame is not received by the original transmitting radio, the unicast frame is not acknowledged and will have to be retransmitted. Additionally, aggregate frames are acknowledged with a Block ACK, and if one of the aggregated frames is corrupted, it will also have to be retransmitted. Any frames that must be retransmitted create extra MAC layer overhead and consume extra airtime in the half-duplex medium.

FIGURE 8 Layer 2 retransmission



Excessive layer 2 retransmissions adversely affect the WLAN in two ways. First, layer 2 retransmissions increase airtime consumption overhead

and therefore decrease throughput. Although other factors may affect throughput, abundant layer 2 retransmissions are usually the culprit.

Second, if application data has to be retransmitted at layer 2, the delivery of application traffic becomes delayed or inconsistent. Applications such as VoIP depend on the timely and consistent delivery of the IP packet. Excessive layer 2 retransmissions usually result in latency and jitter problems for timesensitive applications such as voice and video. When discussing VoIP, people are often confused about the difference between latency and jitter.

Latency Latency is the time it takes to deliver a packet from the source device to the destination device. Ideally, latency should not exceed 50 milliseconds for a VoIP packet. A delay in the delivery (increased latency) of a VoIP packet due to layer 2 retransmissions can result in echo problems.

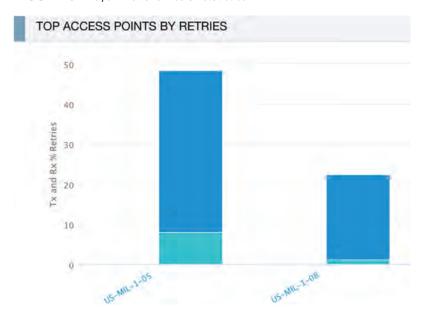
Jitter Jitter is a variation of latency. Jitter measures how much the latency of each packet varies from the average. If all packets travel at exactly the same speed through the network, jitter will be zero. A high variance in the latency (jitter) is a common result of 802.11 layer 2 retransmissions. Jitter will result in choppy audio communications, and constant retransmissions will result in reduced battery life for VoWiFi phones. Although client devices use jitter buffers to compensate for varying delay however, they are usually only effective on delay variations less than 100 milliseconds. Jitter buffers will not compensate for a high percentage of layer 2 retransmission. Jitter variation of less than 5 milliseconds is the ideal goal for VoWiFi.

Most data applications in a Wi-Fi network can handle a layer 2 retransmission rate of up to 10 percent without any noticeable degradation in performance. However, time-sensitive applications such as VoIP require that higher-layer IP packet loss be no greater than 2 percent. Therefore, Voice over Wi-Fi (VoWiFi) networks need to limit layer 2 retransmissions to 5 percent or less to ensure the timely and consistent delivery of VoIP packets. VoWiFi communication usually is restricted to 5 GHz because maintaining a 5 percent layer 2 retry rate in the over-crowded 2.4 GHz band is rarely possible.

How can you measure layer 2 retransmissions? Any good 802.11 protocol analyzer can track layer 2 retry statistics for the entire WLAN. 802.11 protocol analyzers can also track retry statistics for each individual WLAN access point and client station. As shown in Figure 9, layer 2 retry statistics can also usually be centrally monitored using APs across an entire WLAN enterprise from a network management server (NMS). Because the retry statistics are from the perspective of the AP radios, transmit (Tx) statistics indicate a measure of downstream retransmissions from the AP radio, while receive (Rx) statistics indicate a measure of upstream client retransmissions.

Even in pristine RF environments, there will always be some layer 2 retransmissions. The goal should be 10 percent or less and 5 percent or less for voice-grade WLANS. Exceeding a 20 percent retry rate will almost always impact performance.

FIGURE 9 Layer 2 retransmission statistics



Unfortunately, layer 2 retransmissions are a result of many possible problems. Multipath, RF interference, and low signal-to-noise ratio (SNR) are problems that exist at layer 1 yet result in layer 2 retransmissions. Other causes of layer 2 retransmissions include hidden nodes, mismatched power settings, and adjacent channel interference, which are all usually a symptom of improper WLAN design.

RF Interference

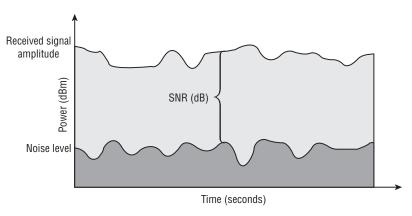
RF interference from a non-802.11 transmitter is the number one cause of layer 2 retransmission. If frames are corrupted due to RF interference, excessive retransmissions will occur, and therefore throughput will be reduced significantly. When layer 2 retransmissions reach excessively high levels intermittently or at different times of the day, the culprit is most likely some sort of interfering device such as a microwave oven. A good WLAN spectrum analyzer will use RF signature files to help you identify the source

of RF inference. To stop the layer 2 retransmissions, locate the interfering device with a spectrum analyzer and remove the interfering device.

Low SNR

Probably the number two most common cause of layer 2 retransmissions is a low SNR. The *signal-to-noise ratio* (*SNR*) is an important value because if the background noise is too close to the received signal or the received signal level is too low, data can be corrupted and layer 2 retransmissions will increase. The SNR is not actually a ratio. It is simply the difference in decibels between the received signal and the background noise (noise floor), as shown in Figure 10. If an 802.11 radio receives a signal of –70 dBm and the noise floor is measured at –95 dBm, the difference between the received signal and the background noise is 25 dB. The SNR is therefore 25 dB.

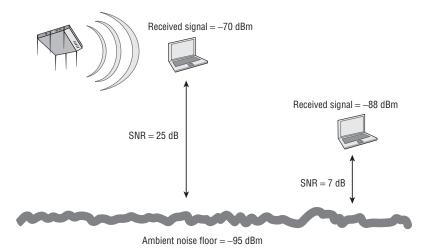
FIGURE 10 Signal-to-noise ratio



Data transmissions can become corrupted with a very low SNR. If the amplitude of the noise floor is too close to the amplitude of the received signal, data corruption will occur and result in layer 2 retransmissions. An SNR of 25 dB or greater is considered good signal quality, and an SNR of 10 dB or lower is considered poor signal quality. To ensure that frames are not corrupted, many vendors recommend a minimum SNR of 20 dB for data WLANs and a minimum SNR of 25 dB for voice WLANs. SNR of lower than 20 dB will result in AP and client radios shifting to a lower *modulation* and coding scheme (MCS) and lower data rates. The lower data rates consume more airtime and degrade performance. Additionally, an SNR of 10 dB or lower will almost guarantee a higher layer 2 retry rate due to data corruption and thus poor performance.

When designing for WLAN coverage, the normal recommended best practice is to provide for a -70 dBm or stronger received signal that is normally well above the noise floor. This will ensure a high SNR. When designing for WLANs with VoWiFi clients, a -65 dBm or stronger signal that is even higher above the noise is recommended. Figure 11 shows a noise floor of -95 dBm. When a client station receives a -70 dBm signal from an access point, the SNR is 25 dB and, therefore, no data corruption results. However, another client receives a weaker -88 dBm signal and a very low SNR of 7 dB. Because the received signal is so close to the noise floor, data corruption will occur and result in layer 2 retransmissions.

FIGURE 11 High and low signal-to-noise ratio

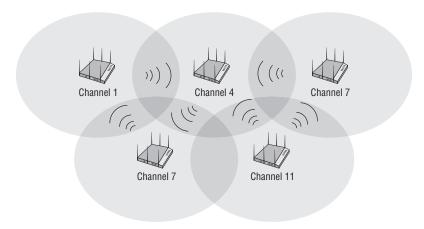


Adjacent Channel Interference

Most Wi-Fi vendors use the term *adjacent channel interference* to refer to degradation of performance resulting from overlapping frequency space that occurs due to an improper channel reuse design. In the WLAN industry, an adjacent channel is considered to be the next or previous numbered channel. For example, channel 3 is adjacent to channel 2. Figure 12 depicts overlapping coverage cells that also have overlapping frequency space, which will result in corrupted data and layer 2 retransmissions. Channels 1 and 4, channels 4 and 7, and channels 7 and 11 all have overlapping frequency space in the 2.4 GHz band. Adjacent channel interference can cause both deferred 802.11 transmissions and corrupted data, which results in layer 2 retries. The performance issues that occur because of adjacent cell interference

normally occur due to a poorly planned 2.4 GHz WLAN. Using a 2.4 GHZ channel reuse pattern of channels 1, 6, and 11 is a standard WLAN design practice that will prevent adjacent cell interference.

FIGURE 12 Adjacent cell interference



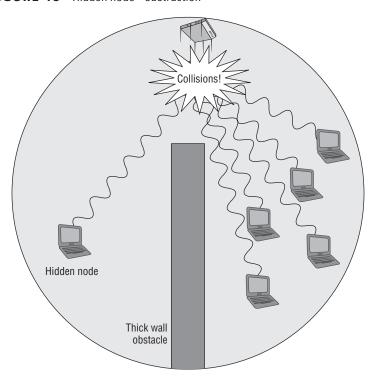
Hidden Node

802.11 radios use a clear channel assessment (CCA) to listen for 802.11 RF transmissions at the Physical layer; the medium must be clear before a station can transmit. The problem with physical carrier sense is that all stations may not be able to hear each other. Remember that the medium is half-duplex and, at any given time, only one radio can be transmitting. What would happen, however, if one client station that was about to transmit performed a CCA but did not hear another station that was already transmitting? If the station that was about to transmit did not detect any RF energy during its CCA, it would transmit. The problem is that you then have two stations transmitting at the same time. The end result is a collision, and the frames will become corrupted. The frames will have to be retransmitted.

The *hidden node* problem occurs when one client station's transmissions are heard by the access point but are not heard by any or all of the other client stations in the basic service set (BSS). The clients would not hear each other and therefore could transmit at the same time. Although the access point would hear both transmissions, because two client radios are transmitting at the same time on the same frequency, the incoming client transmissions would be corrupted.

Figure 13 shows the coverage area of an access point. Note that a thick block wall resides between one client station and all of the other client stations that are associated to the access point. The RF transmissions of the lone station on the other side of the wall cannot be heard by all of the other 802.11 client stations, even though all the stations can hear the AP. That unheard station is the hidden node. What keeps occurring is that every time the hidden node transmits, another station is also transmitting and a collision occurs. The hidden node continues to have collisions with the transmissions from all the other stations that cannot hear it during the clear channel assessment. The collisions continue on a regular basis and so do the layer 2 retransmissions, with the final result being a decrease in throughput. A hidden node can drive retransmission rates above 15 to 20 percent or even higher. Retransmissions, of course, will affect throughput and latency.

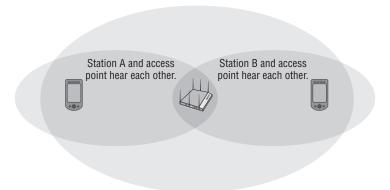
FIGURE 13 Hidden node-obstruction



The hidden node problem may exist for several reasons—for example, poor WLAN design or obstructions such as a newly constructed wall or a newly installed bookcase. A user moving behind some sort of obstruction can cause a hidden node problem. Smartphones and other mobile Wi-Fi devices often become hidden nodes because users take the mobile device into quiet corners or areas where the RF signal of the phone cannot be heard by other client stations. Users with wireless desktops often place their device underneath a metal desk and effectively transform the desktop radio into an unheard hidden node.

The hidden node problem can also occur when two client stations are at opposite ends of an RF coverage cell and they cannot hear each other, as shown in Figure 14. This often happens when the effective coverage cells are too large as a result of the access point's radio transmitting at too much power.

FIGURE 14 Hidden node—large coverage cell



Station A and Station B cannot hear each other.

Another cause of the hidden node problem is distributed antenna systems. Some manufacturers design distributed systems, which are basically made up of a long coaxial cable with multiple antenna elements. Each antenna in the distributed system has its own coverage area. Many companies purchase a distributed antenna system (DAS) for cost-saving purposes. Distributed antenna systems and leaky cable systems are specialty solutions that are sometimes deployed because they can also provide coverage for cellular phone frequencies. The hidden node problem, as shown in Figure 15, will almost always occur if only a single access point is connected to the DAS. If a DAS solution is deployed, multiple APs will still be needed.

Heard signal Unheard transmissions

Access point

FIGURE 15 Hidden node—distributed antenna system

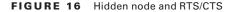
So how do you troubleshoot a hidden node problem? If your end users complain of a degradation of throughput, one possible cause is a hidden node. A protocol analyzer is a useful tool in determining hidden node issues. If the protocol analyzer indicates a higher retransmission rate for the MAC address of one station when compared to the other client stations, chances are a hidden node has been found. Some protocol analyzers even have hidden node alarms based on retransmission thresholds.

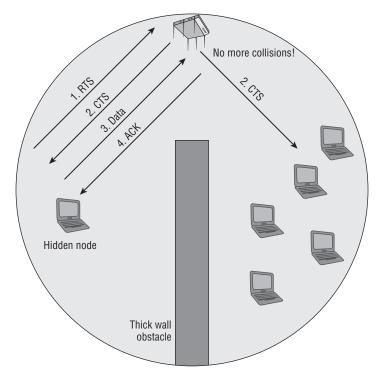
Another way is to use request to send/clear to send (RTS/CTS) to diagnose the problem. If a client device can be configured for RTS/CTS, try lowering the RTS/CTS threshold on a suspected hidden node to about 500 bytes. This level may need to be adjusted depending on the type of traffic being used. For instance, let's say you have deployed a terminal emulation application in a warehouse environment and a hidden node problem exists. In this case, the RTS/CTS threshold should be set for a much lower size, such as 50 bytes. Use a protocol analyzer to determine the appropriate size. RTS/CTS is a method in which client stations can reserve the medium. In Figure 16, you see a hidden node initiating an RTS/CTS exchange.

The stations on the other side of the obstacle may not hear the RTS frame from the hidden node, but they will hear the CTS frame sent by the access point. The stations that hear the CTS frame will reset their NAV for the period of time necessary for the hidden node to transmit the data frame and receive its ACK frame. Implementing RTS/CTS on a hidden node will reserve the medium and force all other stations to pause; thus, the collisions and retransmissions will decrease.

Collisions and retransmissions as a result of a hidden node will cause throughput to decrease. RTS/CTS usually decreases throughput as well. However, if RTS/CTS is implemented on a suspected hidden node, throughput will probably *increase* due to the stoppage of the collisions and

retransmissions. If you implement RTS/CTS on a suspected hidden node and throughput increases, you have confirmed the existence of a hidden node.





Many legacy 802.11 client devices had the ability to adjust RTS/CTS thresholds. In reality, most current client devices cannot be manually configured for RTS/CTS. Therefore, RTS/CTS as a diagnostic tool from the client side usually is not an option. It should be noted that because hidden node problems occur often, WLAN radios may automatically use RTS/CTS to alleviate hidden node problems. Automatic use of RTS/CTS will more likely occur from AP radios as opposed to client-side radios.

RTS/CTS thresholds can always be manually adjusted on access points. A common use of manually adjusted RTS/CTS is *point-to-multipoint* (*PtMP*) bridging. The nonroot bridges in a PtMP scenario will not be able to hear each other because they may be miles apart. RTS/CTS should be

implemented on nonroot PtMP bridges to eliminate collisions caused by hidden node bridges that cannot hear each other.

The following methods can be used to fix a hidden node problem:

Use RTS/CTS. Use either a protocol analyzer or RTS/CTS to diagnose the hidden node problem. RTS/CTS can also be used as an automatic or manual fix to the hidden node problem.

Increase power to all stations. Most client stations have a fixed transmission power output. However, if power output is adjustable on the client side, increasing the transmission power of client stations will increase the transmission range of each station. If the transmission range of all stations is increased, the likelihood of the stations hearing each other also increases. This is usually a bad idea and not a recommended fix because increasing client power can increase co-channel interference.

Remove the obstacles. If it is determined that some sort of obstacle is preventing client stations from hearing each other, simply removing the obstacle will solve the problem. Obviously, you cannot remove a wall, but if a metal desk or file cabinet is the obstacle, it can be moved to resolve the problem.

Move the hidden node station. If one or two stations are in an area where they become unheard, simply moving them within transmission range of the other stations will solve the problem.

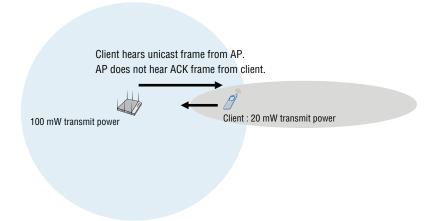
Add another access point. If moving the hidden nodes is not an option, adding another access point in the hidden area to provide coverage will rectify the problem. The best fix for a continuous hidden node problem is to add another AP.

Mismatched Power

Another potential cause of layer 2 retransmissions is mismatched transmit power settings between an access point and a client radio. Communications can break down if a client station's transmit power level is less than the transmit power level of the access point. As a client moves to the outer edges of the coverage cell, the client can "hear" the AP; however, the AP cannot "hear" the client. The good news is that this problem does not occur much in high-density design indoor environments. In recent years there have been significant improvements to access point hardware. Improved receive sensitivity of AP radios has essentially fixed many issues with client and AP mismatched power settings in indoor environments. Problems that occur because of mismatched power settings are more likely to occur outdoors.

As you can see in Figure 17, if an outdoor access point has a transmit power of 100 mW and a client has a transmit power of 20 mW, the client will hear a unicast frame from the AP because the received signal is within the client station's receive sensitivity capabilities. However, when the client sends an ACK frame back to the AP, the amplitude of the client's transmitted signal has dropped well below the receive sensitivity threshold of the AP's radio. The ACK frame is not "heard" by the AP, which then must retransmit the unicast frame. All of the client's transmissions are effectively seen as noise by the AP, and layer 2 retransmissions are the result.

FIGURE 17 Mismatched AP and client power



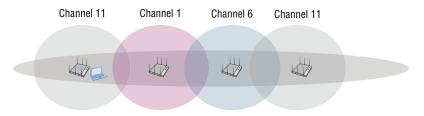
AP/client power problems usually occur because APs are often deployed at full power to increase range. Increasing the power of an access point is the wrong way to increase range. If you want to increase the range for the clients, the best solution is to increase the antenna gain of the access point. Most people do not understand the simple concept of *antenna reciprocity*, which means that antennas amplify received signals just as they amplify transmitted signals. A high-gain antenna on an access point will amplify the AP's transmitted signal and extend the range at which the client is capable of hearing the signal. The AP's high-gain antenna will also amplify the received signal from a distant client station

One way to test whether the mismatched AP/client power problem exists is to listen with a protocol analyzer. An AP/client power problem exists if the frame transmissions of the client station are corrupted when you listen near the access point but are not corrupted when you listen near the client station.

How do you prevent layer 2 retries that are caused by mismatched power settings between the AP and clients? The best solution is to ensure that all of the client transmit power settings match the access point's transmit power. However, significant improvements in AP receive sensitivity have essentially fixed many issues with client and AP mismatch power settings. With that in mind, configuring an access point to transmit at full power is usually not a good idea and may cause this problem as well as many other problems mentioned earlier in this booklet.

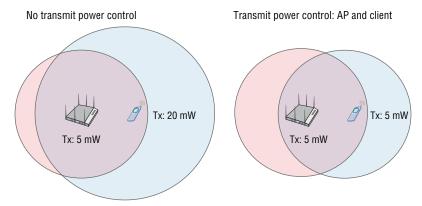
In this section, we have focused on mismatched power settings being a symptom of the too much transmit power from the access point. In reality, the much bigger problem of mismatched power is usually the fact that clients transmit at higher power than access points deployed indoors. The transmit power of many indoor APs may be 10 mW or less due to high density design needs. However, most clients such as smartphones and tablets may transmit at a fixed amplitude of 15 mW or 20 mW. Because clients often transmit at a higher power than the APs and because clients are mobile, the result will be co-channel interference (CCI) as shown in Figure 18. As mentioned earlier in this booklet, CCI results in medium contention overhead which consumes valuable airtime. What most people do not understand about CCI is the fact that clients are the number one cause of CCI. You should understand that CCI is not static and is always changing do to the mobility of client devices.

FIGURE 18 Client-based co-channel interference



An access point with 802.11k capabilities enabled can inform associated clients to use *transmit power control (TPC)* capabilities to change their transmit amplitude dynamically to match the AP's power. Clients that support TPC will adjust their power to match the AP transmit power, as shown in Figure 19. Implementing TPC settings on an AP will greatly reduce co-channel inference caused by clients. It should, however, be noted that legacy clients do not support TPC, and some legacy clients may have connectivity issues if TPC is enabled on the AP.

FIGURE 19 Transmit power control



Multipath

Multipath is an RF propagation phenomenon that results in two or more paths of the same signal arriving at a receiving antenna at the same time or within nanoseconds of each other. Multipath can cause intersymbol interference (ISI), which causes data corruption. Because of the difference in time between the primary signal and the reflected signals, known as the delay spread, the receiver can have problems demodulating the RF signal's information. The delay spread time differential results in corrupted data. If the data is corrupted because of multipath, layer 2 retransmissions will result.

Multipath can be a serious problem when working with legacy 802.11a/b/g equipment. The use of directional antennas will often reduce the number of reflections, and antenna diversity can also be used to compensate for the negative effects of multipath. Multipath is an RF phenomenon that for many years caused destructive effects when older 802.11a/b/g technology was deployed. However, because most WLAN deployments have upgraded to 802.11n or 802.11ac technology, multipath is no longer our enemy. Multipath has a constructive effect with 802.11n/ac transmissions that utilize multiple-input, multiple-output (MIMO) antennas and maximum ratio combining (MRC) signal processing techniques.

Security Troubleshooting

802.11 security defines layer 2 authentication methods and layer 2 dynamic encryption. Therefore, WLAN security problems will usually occur at layer 2 and result in WLAN client connection failures. Many WLAN vendors offer layer 2 diagnostic tools to troubleshoot client device authentication and association. These diagnostic tools may be accessible directly from an AP, or a cloud-based network management system (NMS). Better diagnostic tools may even offer suggested remediation for detected problems. Security and AAA log files from the WLAN hardware and the RADIUS server are also a great place to start when troubleshooting either PSK or 802.1X/EAP authentication problems. Log files may also be gathered from individual WLAN supplicants.

PSK Troubleshooting

Troubleshooting PSK authentication is relatively easy. WLAN vendor diagnostic tools, log files, and a protocol analyzer can all be used to observe the 4-Way Handshake process between a WLAN client and an access point. Let's first take a look at a successful PSK authentication. In Figure 20, you can see the client associate with the AP, and then PSK authentication begins. Because the PSK credentials matched on both the access point and the client, a pairwise master key (PMK) is created to seed the 4-Way Handshake. The 4-Way Handshake process is used to create the dynamically generated unicast encryption key that is unique to the AP radio and the client radio.

Figure 20 shows that the 4-Way Handshake process was successful and that the unicast *pairwise transient key (PTK)* is installed on the AP and the client. The layer 2 negotiations are now complete, and it is time for the client to move on to higher layers. So of course the next step is that the client obtains an IP address via DHCP. If the client does not get an IP address, there is a networking issue and therefore the problem is not a Wi-Fi issue.

Perhaps a Wi-Fi administrator receives a phone call from an end user who cannot get connected using WPA2-Personal. The majority of problems are at the Physical layer; therefore, Wi-Fi Troubleshooting 101 dictates that the end user first enable and disable the Wi-Fi network card. This should ensure that the Wi-Fi NIC drivers are communicating properly with the operating system. If the connectivity problem persists, the problem exists at layer 2. You

can then use diagnostic tools, log files, or a protocol analyzer to observe the failed PSK authentication of the WLAN client.

FIGURE 20 Successful PSK authentication

	o Name Device BSSID Event			
12-A-3BD500	08EA443BD514	Basic	Rx assoc req (rssi 40dB)	
12-A-3BD500	08EA443BD514	Basic	Tx assoc resp <accept> (status 0, pwr 3dBm)</accept>	
12-A-3BD500	08EA443BD514	Info	WPA-PSK auth is starting (at if=wifi0.1)	
12-A-3BD500	08EA443BD514	Info	Sending 1/4 msg of 4-Way Handshake (at if=wifi0:1)	
12-A-3BD500	08EA443BD514	Info	Sending 1/4 msg of 4-Way Handshake (at if=wifi0.1)	
12-A-3BD500	08EA443BD514	Info	Received 2/4 msg of 4-Way Handshake (at if=wifi0.1)	
12-A-3BD500	08EA443BD514	Info	Sending 3/4 msg of 4-Way Handshake (at if=wifi0.1)	
12-A-3BD500	08EA443BD514	Info	Received 4/4 msg of 4-Way Handshake (at thewifi0.1)	
12-A-3BD500	08EA443BD514	Info	PTK is set (at if=wifi0.1)	
12-A-3BD500	08EA443BD514	Basic	Authentication is successfully finished (at if=wifi0.1)	
12-A-3BD500	08EA443BD514	Info	station sent out DHCP DISCOVER message	
12-A-3BD500	08EA443BD514	Info	DHCP server sent out DHCP OFFER message to station	
12-A-3BD500	08EA443BD514	Info	DHCP server sent out DHCP OFFER message to station	
12-A-3BD500	08EA443BD514	Info	station sent out DHCP REQUEST message	
12-A-3BD500	08EA443BD514	Info	DHCP server sent out DHCP ACKNOWLEDGE message to stat	
12-A-3BD500	08EA443BD514	Basic	DHCP session completed for station	
12-A-3BD500	08EA443BD514	Basic	IP 10.5,1.162 assigned for station	

In Figure 21, you can see the client associate and then start PSK authentication. However, the 4-Way Handshake process fails. Notice that only two frames of the 4-Way Handshake complete.

The problem is almost always a mismatch of the PSK credentials. If the PSK credentials do not match, a *pairwise master key (PMK)* seed is not properly created and therefore the 4-Way Handshake fails entirely. The final pairwise transient key (PTK) is never created. A symbiotic relationship exists between authentication and the creation of dynamic encryption keys. If PSK authentication fails, so does the 4-Way Handshake that is used to create the dynamic encryption keys. There is no attempt by the client to get an IP address because the layer 2 process did not complete.

2016-02-22 16:06:48	05-A-764fc0	08EA44764FD4	Info	WPA-PSK auth is starting (at if=wifi0.1)
2016-02-22 16:06:48	05-A-764fc0	08EA44764FD4	Info	Sending 1/4 msg of 4-Way Handshake (at if=wifi0.1)
2016-02-22 16:06:49	05-A-764fc0	08EA44764FD4	Info	Received 2/4 msg of 4-Way Handshake (at if=wifi0.1)
2016-02-22 16:06:52	05-A-764fc0	08EA44764FD4	Info	Sending 1/4 msg of 4-Way Handshake (at it=wifi0.1)
2016-02-22 16:06:52	05-A-764fc0	08EA44764FD4	Info	Received 2/4 msg of 4-Way Handshake (at if=wifi0.1)

FIGURE 21 Unsuccessful PSK authentication

When PSK security is configured, an 8 to 63 character case-sensitive passphrase is entered by the user or administrator. This passphrase is then used to create the PSK. The passphrase could possibly be improperly configured on the access point; however, the majority of the time, the problem is simple: The end user is incorrectly typing in the passphrase. The administrator should make a polite request to the end user to retype the passphrase slowly and carefully, which is a well-known cure for what is known as fat-fingering.

Another possible cause of the failure of PSK authentication could be a mismatch of the chosen encryption methods. An access point might be configured to support only WPA2 (CCMP-AES), which a legacy WPA (TKIP) client does not support. A similar failure of the 4-Way Handshake would occur.

802.1X/EAP Troubleshooting

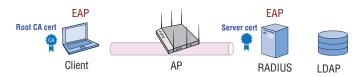
PSK authentication (also known as WPA2-Personal) is simple to trouble-shoot because the authentication method was designed to be uncomplicated. However, troubleshooting the more complex 802.1X/EAP authentication (also known as WPA2-Enterprise) is a bigger challenge because multiple points of failure exist.

802.1X is a port-based access control standard that defines the mechanisms necessary to authenticate and authorize devices to network resources. The 802.1X authorization framework consists of three main components, each with a specific role. These three 802.1X components work together to make sure only properly validated users and devices are authorized to access network resources. The three 802.1X components are known as the supplicant, authenticator, and authentication server. The supplicant is the user or device that is requesting access to network resources. The authentication server's job is to validate the supplicant's credentials. The authenticator is a gateway device that sits in the middle between the supplicant and authentication server, controlling or regulating the supplicant's access to the network.

802.1X/EAP Troubleshooting Zones

In the example shown in Figure 22, the supplicant is a Wi-Fi client, an AP is the authenticator, and an external RADIUS server functions as the authentication server. The RADIUS server can maintain an internal user database or query an external database, such as an LDAP database. Extensible Authentication Protocol (EAP) is used within the 802.1X framework to validate users at layer 2. The supplicant will use an EAP protocol to communicate with the authentication server at layer 2. The Wi-Fi client will not be allowed to communicate at the upper layers of 3–7 until the RADIUS server has validated the supplicant's identity at layer 2.

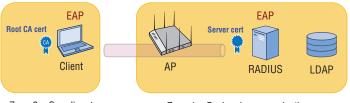
FIGURE 22 802.1X/EAP



The AP blocks all of the supplicant's higher-layer communications until the supplicant is validated. When the supplicant is validated, higher layer communications are allowed through a virtual "controlled port" on the AP (the authenticator). Layer 2 EAP authentication traffic is encapsulated in RADIUS packets between the authenticator and the authentication server. The authenticator and the authentication server also validate each other with a "shared secret."

Better versions of EAP such as EAP-PEAP and EAP-TTLS use "tunneled authentication" to protect the supplicant credentials from offline dictionary attacks. Certificates are used within the EAP process to create an encrypted SSL/TLS tunnel and ensure a secure authentication exchange. As illustrated in Figure 23, a server certificate resides on the RADIUS server and the root CA public certificate must be installed on the supplicant. As mentioned earlier, there are many points of failure in an 802.1X/EAP process. However, as depicted in Figure 23, there are effectively two troubleshooting zones within the 802.1X/EAP framework where failures will occur. Troubleshooting zone 1 consists of the backend communications between the authenticator, the authentication server, and the LDAP database. Troubleshooting zone 2 resides solely on the supplicant device that is requesting access.

FIGURE 23 802.1X/EAP troubleshooting zones



Zone 2 – Supplicant

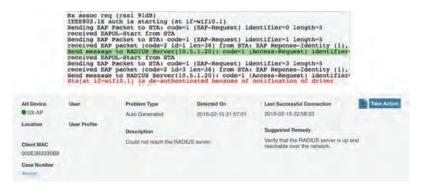
Zone 1 - Backend communications

Zone 1: Backend Communication Problems

Zone 1 should always be investigated first. If an AP and a RADIUS server cannot communicate with each other, the entire authentication process will fail. If the RADIUS server and the LDAP database cannot communicate, the entire authentication process will also fail.

Figure 24 shows a capture of a supplicant (Wi-Fi client) trying to contact a RADIUS server. The authenticator forwards the request to the RADIUS server, but the RADIUS server never responds. The AP (authenticator) then sends a deauthentication frame to the Wi-Fi client because the process failed. This is an indication that there is a backend communication problem in the first troubleshooting zone.

FIGURE 24 The RADIUS server does not respond.

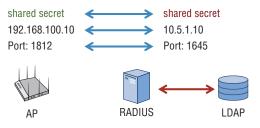


As shown in Figure 25, if the RADIUS server never responds to the supplicant, there are four possible points of failure in the first troubleshooting zone:

- Shared secret mismatch
- Incorrect IP settings on the AP or the RADIUS server

- Authentication port mismatch
- LDAP query failure

FIGURE 25 Points of failure -802.1X/EAP troubleshooting zone 1



- · Shared secret mismatch
- Incorrect IP settings on AP or RADIUS server
- Authentication port mismatch (default is 1812)
- LDAP communications error

The first three possible points of failure are between the authenticator and the RADIUS server. The authenticator and the authentication server validate each other with a *shared secret*. The most common failure in RADIUS communications is that the shared secret has been typed in wrong on either the RADIUS server or the AP functioning as the authenticator.

The second most common failure in RADIUS communications is simply misconfigured IP networking settings. The AP must know the correct IP address of the RADIUS server. Likewise the RADIUS server must be configured with the IP addresses of any APs functioning as authenticators. Incorrect IP settings will result in miscommunications.

The third point of failure between an authenticator and an authentication server is a mismatch of RADIUS authentication ports. UDP ports 1812 and 1813 are defined as the industry standard ports used for RADIUS authentication and accounting. However, some older RADIUS servers may be using UDP ports 1645 and 1646. UDP ports 1645 and 1646 are rarely used anymore but do occasionally show up on older RADIUS servers. Although not a common point of failure, if the authentication ports do not match between a RADIUS server and the AP, the authentication process will fail.

The final point of failure on the backside is a failure of the LDAP query between a RADIUS server and the LDAP database. A standard domain account can be used for LDAP queries; however, if the account has expired or if there is a networking issue between the RADIUS server and the LDAP server, the entire 802.1X/EAP authentication process will fail.

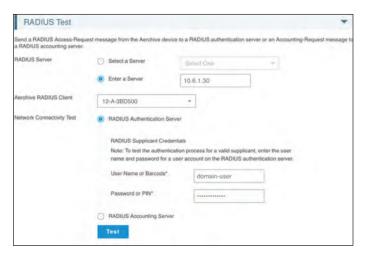


Real World Scenario

What Tools Can Be Used to Troubleshoot 802.1X/EAP Backend Communications?

The good news is that multiple troubleshooting resources are available to troubleshoot zone 1. Several WLAN vendors offer built-in diagnostic tools to test the communications between an authenticator and a RADIUS server as well as LDAP communications. Aerohive access points function as the 802.1X authenticator that communicates with the RADIUS server. As shown in Figure 26, a standard domain account and password can be used to test the RADIUS and EAP communications.

FIGURE 26 802.1X/EAP backend diagnostic tool



Several software utilities are also available to test backend 802.1X/EAP communications. EAPTest is a commercial test utility available for the macOS. EAPTest is available from the Mac App Store. More information can be found at www.ermitacode.com/eaptest.html. RADLogin is a free test utility for the Windows and Linux platforms. More information can be found at www.iea-software.com/products/radlogin4.cfm. RADIUS server and LDAP database logs are also great resources for troubleshooting 802.1X/EAP backend communication problems. In worst-case scenarios, a wired protocol analyzer may be needed to capture RADIUS packets. Many of these test tools can also be used to troubleshoot issues with RADIUS attributes, which can be leveraged during 802.1X/EAP authentication for *rolebased access control (RBAC)*.

Zone 2: Supplicant Certificate Problems

If all backend communications between the authenticator and the RADIUS server are functioning properly, then the 802.1X/EAP troubleshooting focus should now be redirected to zone 2. In simpler words, the culprit is the client (supplicant). Problems with the supplicant usually either revolve around certificate issues or client credential issues. Let's take a look at Figure 27. Note that the RADIUS server is responding and therefore verifying that the backend communications are good. Also notice an SSL tunnel negotiation starts and finishes successfully. This 802.1X/EAP diagnostic log confirms that the certificate exchange was successful and that an SSL/TLS tunnel was successfully created to protect the supplicant credentials.

FIGURE 27 Successful SSL/TLS tunnel creation

Send message to RADIUS Server(10.5.1.129); code=1 (Access-Request) identifier=RADIUS: SSL negotiation, send server certificate and other message Receive message from RADIUS Server: code=11 (Access-Challenge) identifier=109 Sending EAP Packet to STA: code=1 (EAP-Request) identifier=3 length=280 received EAP packet (code=2 id=3 len=208) from STA: EAP Reponse-PEAP (25) Send message to RADIUS Server(10.5.1.129); code=1 (Access-Request) identifier=RADIUS: SSL connection established
Receive message from RADIUS Server: code=11 (Access-Challenge) identifier=110 Sending EAP Packet to STA: code=1 (EAP-Request) identifier=4 length=65 received EAP packet (code=2 id=4 len=6) from STA: EAP Reponse-PEAP (25) Send message to RADIUS Server(10.5.1.129); code=1 (Access-Request) identifier=RADIUS: SSL negotiation is finished successfully
Receive message from RADIUS Server: code=11 (Access-Challenge) identifier=111 Sending EAP Packet to STA: code=1 (EAP-Request) identifier=5 length=43 received EAP packet (code=2 id=5 len=59) from STA: EAP Reponse-PEAP (25) Send message to RADIUS Server(10.5.1.129); code=1 (Access-Request) identifier=RADIUS: PEAP inner tunneled conversion

Figure 28 displays an 802.1X/EAP diagnostic log where you can see the SSL negotiation begin and the server certificate sent from the RADIUS server to the supplicant. However, the SSL/TLS tunnel is never created, and EAP authentication fails. If the SSL/TLS tunnel cannot be established, this is an indication that there is some sort of certificate problem.

FIGURE 28 Unsuccessful SSL/TLS tunnel creation

Rx assoc req (rssi 95dB)
IEEE802.1X auth is starting (at if=wifi0.1)
Sending EAP Packet to STA: code=1 (EAP-Request) identifier=0 length=5
received EAP packet (code=2 id=0 len=16) from STA: EAP Reponse-Identity (1),
Send message to RADIUS Server(10.5.1.129): code=1 (Access-Request) identifier
RADIUS: EAF start with type peap
Receive message from RADIUS Server: code=11 (Access-Challenge) identifier=50
Sending EAP Packet to STA: code=1 (EAP-Request) identifier=1 length=6
received EAP packet (code=2 id=1 len=105) from STA: EAP Reponse-FEAP (25)
Send message to RADIUS Server(10.5.1.129): code=1 (Access-Request) identifier=8
RADIUS: SSL negotiation, receive client hello message
Receive message from RADIUS Server: code=11 (Access-Challenge) identifier=51
Sending EAP Packet to STA: code=1 (EAP-Request) identifier=2 length=1024
received EAP packet (code=2 id=2 len=6) from STA: EAP Reponse-FEAP (25)
Send message to RADIUS Server(10.5.1.129): code=1 (Access-Request) identifier
RADIUS: SSL negotiation, send server certificate and other message
Receive message from RADIUS Server: code=11 (Access-Challenge) identifier=52
Sending EAP Packet to STA: code=1 (EAP-Request) identifier=3 length=280
received EAP packet (code=2 id=3 len=6) from STA: EAP Reponse-FEAP (25)
Send message to RADIUS Server(10.5.1.129): code=1 (Access-Request) identifier
RADIUS: SSL negotiation, send server certificate and other message
Receive message from RADIUS Server: code=11 (Access-Challenge) identifier=53
Sending EAP Packet to STA: code=1 (EAP-Request) identifier=53
Sending EAP Packet to STA: code=1 (EAP-Request) identifier=653
Sending EAP Packet to STA: code=1 (EAP-Request) identifier=654
Sending EAP Packet to STA: code=1 (EAP-Request) identifier=654

You can usually verify that there is a certificate problem by editing the supplicant client software settings and temporarily disabling the validation of the server certificate, as shown in Figure 29. If EAP authentication is successful after you temporarily disable the validation of the server certificate, then you have confirmed there is a problem with the implementation of the certificates within the 802.1X framework. Please note that this is not a fix but an easy way to verify that some sort of certificate issue exists.

FIGURE 29 Server certificate validation



A whole range of certificate problems could be causing the SSL/TLS tunnel not to be successfully created. The most common certificate issues are as follows:

- The root CA certificate is installed in the incorrect certificate store.
- The incorrect root certificate is chosen.
- The server certificate has expired.
- The root CA certificate has expired.
- The supplicant clock settings are incorrect.

The root CA certificate needs to be installed in the Trusted Root Certificate Authorities store of the supplicant device. A common mistake is to install the root CA certificate in the default location, which is typically the personal store of a Windows machine. Another common mistake is to select the incorrect root CA certificate with the supplicant configuration. The SSL/

TLS tunnel will fail because the incorrect root CA certificate will not be able to validate the server certificate. Digital certificates are also time-based, and a common problem is that the server certificate has expired. Although not as common, the root CA certificate can also have expired. The clock settings on the supplicant may be incorrect and might possibly predate the creation of either certificate.

Because of all the possible points of failure involving certificates, troubleshooting 802.1X/EAP certificate problems in zone 2 can be difficult. Additionally, there are more potential problems with certificates. The server certificate configuration may be incorrect on the RADIUS server. In other words, the certificate problem exists back in troubleshooting zone 1. What if EAP-TLS is the deployed authentication protocol? EAP-TLS requires the provisioning of client-side certificates in addition to server certificates. Client certificates add an additional layer of possible certificate troubleshooting on the supplicant as well as within the private PKI infrastructure that has been deployed.

There is one final complication that might result in the failure of tunneled authentication. The chosen layer 2 EAP protocols must match on both the supplicant and the authentication server. For example, the authentication will fail if PEAPv0 (EAP-MSCHAPv2) is selected on the supplicant while PEAPv1 (EAP-GTC) is configured on the RADIUS sever. Although the SSL/TLS tunnel might still be created, the inner tunnel authentication protocol does not match and authentication will fail. Although it is possible for multiple flavors of EAP to operate simultaneously over the same 802.1X framework, the EAP protocols must match on both the supplicant and the authentication server.

Zone 2: Supplicant Credential Problems

If you can verify that you do not have any certificate issues and the SSL/TLS tunnel is indeed established, the supplicant problems are credential failures. Figure 30 displays an 802.1X/EAP diagnostic log where the RADIUS server is rejecting the supplicant credentials. The following supplicant credential problems are possible:

- Expired password or user account
- Wrong password
- User account does not exist in LDAP.
- Machine account has not been joined to the Windows domain.

FIGURE 30 RADIUS server rejects supplicant credentials

```
RADIUS: SSL connection established
Receive message from RADIUS Server: code=11 (Access-Challenge) identifier=127 length=123
Sending EAP Packet to STA: code=1 (EAP-Request) identifier=5 length=65
received EAP packet (code=2 id=5 lene6) from STA: EAP Reponse-PEAP (25)
Send message to RADIUS Server(10.5.1.129): code=1 (Access-Request) identifier=128 length=176
RADIUS: SSL megotiation is finished successfully
Receive message from RADIUS Server: code=11 (Access-Challenge) identifier=128 length=101
Sending EAP Packet to STA: code=1 (EAP-Request) identifier=6 length=43
received EAP packet (code=2 id=6 len=43) from STA: EAP Reponse-PEAP (25)
Send message to RADIUS Server: (10.5.1.129): code=1 (Access-Request) identifier=129 length=123
RADIUS: PEAP inner tunneled conversion
Receive message from RADIUS Server: code=1 (Access-Challenge) identifier=129 length=117
Sending EAP Packet to STA: code=1 (EAP-Request) identifier=7 length=59
received EAP packet (code=2 id=7 len=91) from STA: EAP Reponse-PEAP (25)
Send message to RADIUS Server: (10.5.1.129): code=1 (Access-Request) identifier=130 length=261
RADIUS: PEAP Tunneled authentication was rejected. NTLM auth failed for Logon failure (0xc00
Receive message from RADIUS Server: code=1 (Access-Challenge) identifier=130 length=101
Sending EAP Packet to STA: code=1 (EAP-Request) identifier=8 length=43
received EAP packet (code=2 id=8 len=43) from STA: EAP Reponse-PEAP (25)
Send message to RADIUS Server: (10.5.1.129): code=1 (Access-Request) identifier=131 length=13
RADIUS: rejected user 'user' through the NAS at 10.5.1.129.
Authentication is terminated (at 12-wif10.1) because it is rejected by RADIUS server
Sending EAP Packet to STA: code=4 (EAP-Failure) identifier=8 length=4
Sta(at 16-wif10.1) is de-authenticated because of notification of driver
```

If the user credentials do not exist in the LDAP database or the credentials have expired, authentication will fail. Unless single sign-on capabilities have been implemented on the supplicant, there is always the possibility that the domain user password can be incorrectly typed by the end user.

Another common error is that the Wi-Fi supplicant has been improperly configured for machine authentication and the RADIUS server has only been configured for user authentication. In Figure 31 we see a diagnostic log that clearly shows the machine credentials being sent to the RADIUS server and not the user credentials. The RADIUS server was expecting a user account and therefore rejected the machine credentials because no machine accounts had been set up for validation. In the case of Windows, the machine credentials are based on a *System Identifier (SID)* value that is stored on a Windows domain computer after being joined to a Windows domain with Active Directory.

FIGURE 31 Machine authentication failure

```
Send message to RADIUS server[10.5.1.129]: code-1 (Access-Request) identifier=151 length=203, RADIUS: SSL negotiation, send server certificate and other message Receive message from RADIUS Server: code=11 (Access-Challenge) identifier=151 length=340 Sending EAP Packet to STA: code=1 (EAP-Request) identifier=4 length=280 received EAP packet (code=2 id=4 len=17) from STA: EAP Reponse-PEAP (25) Send message to RADIUS Server[10.5.1.129]: code=1 (Access-Request) identifier=152 length=214, RADIUS: rejected user 'host/TRAINING-PC16.ah-lab.local' through the NAS at 10.5.1.129. Authentication is terminated (at if=wifi0.1) because it is rejected by RADIUS server Sending EAP Packet to STA: code=4 (EAP-Failure) identifier=4 length=4 Sta(at if=wifi0.1) is de-authenticated because of notification of driver
```

Of course, a WLAN administrator can always verify that all is well with an 802.1X/EAP client session. Always remember that a byproduct of the EAP process is the generation of the pairwise master key (PMK) that seeds the 4-Way Handshake exchange. Figure 32 shows the EAP process completing;

the pairwise master key (PMK) is sent to the AP from the RADIUS server. The 4-Way Handshake process then begins to dynamically generate the pairwise transient key (PTK) that is unique between the radios of the AP and the client device. When the 4-Way Handshake completes, the encryption keys are installed and the layer 2 connection is completed. The virtual controlled port on the authenticator opens up for this Wi-Fi client. The supplicant can now proceed to higher layers and get an IP address. If the client does not get an IP address, there is a networking issue and therefore the problem is not a Wi-Fi issue.

FIGURE 32 4-Way Handshake

```
Receive message from RADIUS Server: code=2 (Access-Accept) identifier=125 PMK is got from RADIUS server (at if=wifi0.1) Sending EAP Packet to STA: code=3 (EAP-Success) identifier=5 length=4 Sending 1/4 mag of 4-Way Handshake (at if=wifi0.1) Received 2/4 mag of 4-Way Handshake (at if=wifi0.1) Received 4/4 mag of 4-Way Handshake (at if=wifi0.1) Received 4/4 mag of 4-Way Handshake (at if=wifi0.1) PTK is set (at if=wifi0.1) Authentication is successfully finished (at if=wifi0.1) IP 10.5.10 assigned for station station sent out DHCP RCQUEST message DHCP server sent out DHCP ACRONMLEDGE message to station DHCP session completed for station
```

One final consideration when troubleshooting 802.1X/EAP is RADIUS attributes. RADIUS attributes can be leveraged during 802.1X/EAP authentication for role-based access control, providing custom settings for different groups of users or devices. For example, different groups of users may be assigned to different VLANs even though they are connected to the same 802.1X/EAP SSID. If the RADIUS attribute configuration does not match on the authenticator and the RADUS server, users might be assigned to default role or VLAN assignments. In worst-case scenarios, a RADIUS attribute mismatch might result in authentication failure.

VPN Troubleshooting

VPNs are rarely used anymore as the primary method of security for WLANs. Occasionally, a VPN may be used to provide data privacy across a point-to-point 802.11 wireless bridge link. IPsec VPNs are still commonly used to connect remote branch offices with corporate offices across WAN links. Although a site-to-site VPN link is not necessarily a WLAN security solution, the wireless user traffic that originated at the remote location may be required to traverse through a VPN tunnel. Most WLAN vendors also offer VPN capabilities within their solution portfolio. For example, a WLAN may offer a VPN solution where user traffic is tunneled from a remote AP or a WLAN branch router to a VPN server gateway. Third-party VPN overlay solutions are often also used.

The creation of an IPsec tunnel involves two phases, called *Internet Key Exchange (IKE)* phases:

IKE Phase 1

The two VPN endpoints authenticate one another and negotiate keying material. The result is an encrypted tunnel used by Phase 2 for negotiating the *Encapsulating Security Payload (ESP)* security associations.

IKE Phase 2

The two VPN endpoints use the secure tunnel created in Phase 1 to negotiate ESP *security associations (SAs)*. The ESP SAs are used to encrypt user traffic that traverses between the endpoints.

The good news is that any quality VPN solution offers diagnostic tools and commands to troubleshoot both IKE phases. Some of the common problems that can occur if IKE Phase 1 fails are as follows:

- Certificate problems
- Incorrect networking settings
- Incorrect NAT settings on the external firewall

In Figure 33 you see the results of an IKE Phase 1 diagnostic command executed on a VPN server. IPsec uses digital certificates during Phase 1. If IKE Phase 1 fails due to a certificate problem, ensure that you have the correct certificates installed properly on the VPN endpoints. Also remember that certificates are time-based. Very often, a certificate problem during IKE Phase 1 is simply an incorrect clock setting on either VPN endpoint.

FIGURE 33 IPsec Phase 1—certificate failure

```
Show MKE Event -05.A-051ec0

2010-04-29 11:23:43:Peer failed phase 1 authentication (certificate prok 2010-04-29 11:24:34:Peer not responding(10.5.1.160[4500]->1.1.1.5[4500]) 2010-04-29 11:24:34:Phase 1 deleted(10.5.1.160[4500]->1.1.1.5[4500]) 2010-04-29 11:24:35:Peer failed phase 1 authentication (certificate prok 2010-04-29 11:25:25:Peer not responding(10.5.1.160[4500]->1.1.1.5[4500]) 2010-04-29 11:25:25:Phase 1 deleted(10.5.1.160[4500]->1.1.1.5[4500]) 2010-04-29 11:25:25:Phase 1 deleted(10.5.1.160[4500]->1.1.1.5[4500]) 2010-04-29 11:25:26:Peer failed phase 1 authentication (certificate prok 2010-04-29 11:26:16:Peer not responding(10.5.1.160[4500]->1.1.1.5[4500]) 2010-04-29 11:26:17:Peer failed phase 1 authentication (certificate prok 2010-04-29 11:27:07:Peer not responding(10.5.1.160[4500]->1.1.1.5[4500]) 2010-04-29 11:27:07:Peer not responding(10.5.1.160[4500]->1.1.1.5[4500]) 2010-04-29 11:27:08:Peer failed phase 1 authentication (certificate prok 2010-04-29 11:27:58:Peer not responding(10.5.1.160[4500]->1.1.1.5[4500]) 2010-04-29 11:27:59:Peer failed phase 1 authentication (certificate prok 2010-04-29 11:28:49:Peer not responding(10.5.1.160[4500]->1.1.1.5[4500]) 2010-04-29 11:29:39:Peer failed phase 1 authentication (certificate prok 2010-04-29 11:29:39:Peer failed phase 1 authentication (certificate prok 2010-04-29 11:29:39:Peer failed phase 1 authentication (ce
```

In Figure 34 you see the results of an IKE Phase 1 diagnostic command executed on a VPN server that indicates a possible networking error due to incorrect configuration. IPsec uses private IP addresses for tunnel communications and also uses external IP addresses, which are normally the public IP address of a firewall. If an IKE Phase 1 failure occurs as shown in Figure 34, check the internal and external IP settings on the VPN devices. If an external firewall is being used, also check the *Network Address Translation (NAT)* settings. Another common networking problem that causes VPNs to fail is that needed firewall ports are blocked. Ensure that the following ports are open on any firewall that the VPN tunnel may traverse:

- UDP 500 (IPsec)
- UDP 4500 (NAT Transversal)

FIGURE 34 IPsec Phase 1-networking failure

show IKE Event	-02-A-700a40)
2013-01-12	12:52:56:Phase	1	started(10.5.2.2[500]->1.2.2.3[500])	
2013-01-12	12:53:07:Phase	1	deleted(10.5.2.2(500)->1.2.2.3(500))	
2013-01-12	12:53:10:Phase	1	started(10.5.2.2(5001->1.2.2.3(5001)	
2013-01-12	12:53:20:Phase	1	deleted(10.5.2.2[500]->1.2.2.3[500])	
2013-01-12	12:53:23:Phase	1	started(10.5.2.2[500]->1.2.2.3[500])	
2013-01-12	12:53:34:Phase	1	deleted(10.5.2.2[500]->1.2.2.3[500])	
2013-01-12	12:53:37:Phase	1	started(10.5.2.2[500]->1.2.2.3[500])	
2013-01-12	12:53:48:Phase	1	deleted(10.5.2.2[500]->1.2.2.3[500])	
2013-01-12	12:53:51:Phase	1	started(10.5.2.2[500]->1.2.2.3[500])	
2013-01-12	12:54:02:Phase	1	deleted(10.5.2.2[500]->1.2.2.3[500])	
2013-01-12	12:54:05:Phase	1	started(10.5.2.2[500]->1.2.2.3[500])	
2013-01-12	12:54:16:Phase	1	deleted(10.5.2.2(500)->1.2.2.3(500))	
2013-01-12	12:54:19:Phase	1	started(10.5.2.2[500]->1.2.2.3[500])	
2013-01-12	12:54:30:Phase	1	deleted(10.5.2.2[500]->1.2.2.3[500])	
2013-01-12	12:54:33:Phase	1	started(10.5.2.2[500]->1.2.2.3[500])	
2013-01-12	12:54:44:Phase	1	deleted(10.5.2.2[500]->1.2.2.3[500])	
2013-01-12	12:54:47:Phase	1	started(10.5.2.2[500]->1.2.2.3[500])	
2013-01-12	12:54:58:Phase	1	deleted(10.5.2.2[500]->1.2.2.3[500])	
2013-01-12	12:55:01:Phase	1	started(10.5.2.2[500]->1.2.2.3[500])	
2013-01-12	12:55:12:Phase	1	deleted(10.5.2.2[500]->1.2.2.3[500])	
2013-01-12	12:55:15:Phase	1	started(10.5.2.2[500]->1.2.2.3[500])	

If you can confirm that IKE Phase 1 is successful yet the VPN is still failing, then IKE Phase 2 is the likely culprit. The following list includes some of the common problems if IKE Phase 2 fails:

- Mismatched transform sets between the client and server (encryption algorithm, hash algorithm, etc.)
- Mixing different vendor solutions

In Figure 35 you see the successful results of an IKE Phase 2 diagnostic command executed on a VPN server. If this command had indicated a failure, be sure to check both encryption and hash settings on the VPN endpoints. Check other IPsec settings such as *tunnel mode*. You will need to verify that all settings match on both ends. IKE Phase 2 problems often occur when different VPN vendors are used on opposite sides of the intended

VPN tunnel. Although IPsec is a standards-based suite of protocols, mixing different VPN vendor solutions often results in more troubleshooting.

FIGURE 35 IPsec Phase 2-success

```
Show Psec SA-02-A-066600

SA(Security Association) information as following:

IPsec Security Association Information:
10.5.1.165 [4500] 1.2.1.2 [4500]

tunnel-id:

esp-udp modestunnel spi=158846310(0x0977cd66) regid=0(0x000000000 Encryption aes-cbc
Authentication: hmac-shal seq=0x00000000 relay=4 flags=0x20000000 state=mature created: Jul 12 12:48:37 2011 current: Jul 12 12:51:31 2011 diff: 174(s) hard: 3600(s) soft: 2880(s) last: Jul 12 12:48:37 2011 hard: 0(s) soft: 0(s) current: 2880(bytes) hard: 0(bytes) soft: 0(bytes) failed: 0(pkts) replay (0pkts) soft: 0(bytes) failed: 0(pkts) replay: 0(pkts) replay window: 0(pkts) sadb seq=1 pid=944 refcnt=0

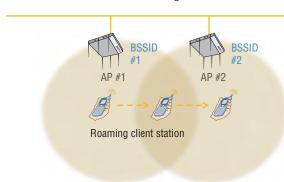
1.2.1.2 [4500] 10.5.1.165 [4500]

tunnel-id: 2 esp-udp mode=tunnel spi=218804365(0x0d0ab08d) reqid=0(0x000000000 Encryption: aes-cbc Authentication: hmsc-shal seq=0x00000000 relay=4 flags=0x20000000 state=mature created: Jul 12 12:48:37 2011 current: Jul 12 12:51:31 2011 diff: 174(s) bard: 3600(s)
```

Roaming Troubleshooting

Mobility is the whole point behind wireless network access. 802.11 clients need the ability to seamlessly roam between access points without any interruption of service or degradation of performance. As shown in Figure 36, seamless roaming has become even more important in recent years because of the proliferation of handheld personal Wi-Fi devices such as smartphones and tablets.

FIGURE 36 Seamless roaming



The most common roaming problems are a result of either bad client drivers or bad WLAN design. The very common *sticky client problem* is when client stations stay connected to their original AP and do not roam to a new AP of closer vicinity and stronger signal. The sticky client problem is often a result of APs in close physical vicinity with transmit power levels that are too high. The sticky client problem and other roaming performance issues can usually be avoided with proper WLAN design and site surveys. Good roaming design entails defining proper primary coverage and secondary coverage from the client perspective.

Client stations, and not the access point, make the decision on whether or not to roam between access points. Some vendors may involve the access point in the roaming decision, but ultimately, the client station initiates the roaming process with a reassociation request frame. The method by which a client station decides to roam depends on unique thresholds determined by the manufacturer of the 802.11 client radio. Roaming thresholds are usually defined by RSSI and SNR; however, other variables such as error rates and retransmissions may also have a part in the roaming decision. Client stations that support 802.11k may obtain neighbor reports from 802.11k-compliant APs, which provide the client stations with additional input so the client stations can make better roaming decisions. Support for 802.11k is becoming more and more important in today's complex RF environments.

Roaming problems will occur if there is not enough duplicate secondary coverage. No secondary coverage will effectively create a roaming dead zone, and connectivity might even temporarily be lost. On the flip side, too much secondary coverage will also cause roaming problems. For example, a client station may stay associated with its original AP and not connect to a second access point even though the station is directly underneath the second access point. As previously mentioned, this is commonly referred to as the sticky client problem. Too many potential APs heard by a client may also result in a situation in which the client device is constantly switching back and forth between the two or more APs on different channels. If a client station can also hear dozens of APs on the same channel with very strong signals, a degradation in performance will occur due to medium contention overhead.

Roaming performance also has a direct relationship to WLAN security. Every time a client station roams, new encryption keys must be generated between the AP and the client station radios via the 4-Way Handshake. When using 802.1X/EAP security, roaming can be especially troublesome for VoWiFi and other time-sensitive applications. Due to the multiple frame exchanges between the authentication server and the supplicant, an

802.1X/EAP authentication can take 700 milliseconds (ms) or longer for the client to authenticate. VoWiFi requires a handoff of 150 ms or less to avoid a degradation of the quality of the call, or even worse, a loss of connection. Therefore, faster, secure roaming handoffs are required.

Changes in the WLAN environment can also cause roaming headaches. RF interference will always affect the performance of a wireless network and can make roaming problematic as well. Very often new construction in a building will affect the coverage of a WLAN and create new dead zones. If the physical environment where the WLAN is deployed changes, the coverage design might have to change as well. It is always a good idea to periodically conduct a validation survey to monitor changes in coverage patterns.

Troubleshooting roaming by using a protocol analyzer is tricky because the reassociation roaming exchanges occur on multiple channels. For example, in order to troubleshoot a client roaming between channels 1, 6, and 11, you would need three separate protocol analyzers on three separate laptops that would produce three separate frame captures. CACE Technologies offers a product called AirPcap that is a USB 802.11 radio. As shown in Figure 37, three AirPcap USB radios can be configured to capture frames on channels 1, 6, and 11 simultaneously. All three radios are connected to a USB hub and save the frame captures of all three channels into a single time-stamped capture file. The AirPcap solution allows for multichannel monitoring with a single protocol analyzer. Other WLAN analyzer vendors also offer multichannel monitoring capabilities for both the 2.4 GHz and 5 GHz frequency bands. The roaming history of a WLAN client can also be collected from AP log files and visualized in WLAN network management solutions.

FIGURE 37 Multichannel monitoring and analysis



Certified Wireless Security Professional Study Guide (Coleman, Harkins and Westcott Sybex, 2016) devotes an entire chapter to fast secure foaming such as opportunistic key caching (OKC) and fast BSS transition (FT). Both OKC and FT produce roaming handoffs of closer to 50 ms even when 802.1X/EAP is the chosen security solution. Both OKC and FT use key distribution mechanisms so that roaming clients do not have to reauthenticate every time they roam. OKC is now considered a legacy method of fast secure roaming. The FT roaming mechanisms defined in both 802.11r and Voice Enterprise are considered the standard. Many WLAN enterprise vendor APs are now certified for Voice Enterprise by the Wi-Fi Alliance. Please note that any client devices that were manufactured before 2012 simply will not support 802.11r/k/v operations. Although the bulk of legacy client devices do not support Voice Enterprise capabilities, client-side support is growing and more commonplace.

Most security-related roaming problems are based on the fact that many clients simply do not support either OKC or fast BSS transition (FT). Client-side support for any device that will be using voice applications and 802.1X/EAP is critical. Proper planning and verification of client-side and AP support for OKC or FT will be necessary. Figure 38 shows the results of a diagnostic command that displays the roaming cache of an access point. This type of diagnostic command can verify if PMKs are being forwarded between access points. In this situation, Voice Enterprise is enabled on the AP and supported on the client radio. You can verify the MAC address of the supplicant and the authenticator as well as the PMK information. Always remember that the supplicant must also support Voice Enterprise; otherwise, the suppliant will reauthenticate every time the client roams.

As mentioned previously, enabling Voice Enterprise mechanisms on an access point may actually create connectivity problems for legacy clients. When FT is configured on an access point, the AP will broadcast management frames with new information elements. For example, the *mobility domain information element (MDIE)* will be in all beacon and probe response frames. Unfortunately, the drivers of some older legacy client radios may not be able to process the new information in these management frames. The result is that legacy clients may have connectivity problems when an AP is configured for FT. Always test the legacy client population when configuring APs for fast BSS transition. If connectivity problems arise, consider using a separate SSID solely for fast BSS transition devices. However, please remember that every SSID consumes airtime due to the layer 2 management

overhead. Additionally, as more devices begin to support Voice Enterprise capabilities, upgrade your client devices.

FIGURE 38 Roaming cache

sh roam cache mac b844:d90e:006e Supplicant Address(SPA): b844:d90e:006e PMK(1st 2 bytes): n/a PMKID(1st 2 bytes): n/a Session time: -1 seconds (-1 means infinite) PMK Time left in cache: 3581 PMK age: 1040 Roaming cache update interval: 60 last time logout: 1221 seconds ago Authenticator Address: MAC=9c5d:122e:c124, IP=172.16.255.93 Roaming entry is got from neighbor AP: 9c5d:122e:c124 PMK is got(Flag): Locally Station IP address: 172.16.255.90 (from DHCP) Station hostname: Davids-IPhone Station default gateway: 172.16.255.1 Station DNS server: 172.16.255.1 Station DHCP lease time: 85349 seconds Hops: 0 WPA key mgmt: 64 BOKH: 9c5d:1263:6464 ROKH IP: 172.16.255.94 PMKR0 Name: 19D2*

Because 802.11 wireless networks are usually integrated into preexisting wired topologies, crossing layer 3 boundaries is often a necessity, especially in large enterprise deployments. The only way to maintain upper-layer communications when crossing layer 3 subnets is to provide a *layer 3 roaming* solution. When clients roam to a new subnet, a generic routing encapsulation (GRE tunnel) must be created to the original subnet so that the WLAN client can maintain its original IP address. As shown in Figure 39, the major WLAN vendors offer diagnostic tools and commands to verify that layer 3 roaming tunnels are being successfully created.

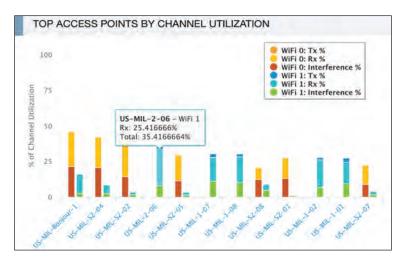
FIGURE 39 GRE tunnel

```
Show GRE Tunnel
                                                                                       X
Tunnel table:
T=Type; Z=Zone; F=Flag; PN=policy numbers;
Age Out-idle time of the tunnel since last receive packet
TXs=TX packets; TXE=TX errors; RXs=RX packets; RXE=RX errors;
Type: G=General route encapsulation; O=Other tunnel;
Zone: A=Access: B=Backhaul;
Total entries: 1
ID T Z F PN Age Out Src IP
TXs
      TXE RXs RXE
  GAPI 0
                    172.28.242.142
                                                          172.28.242.141
121 0 54
                    0
```

Channel Utilization

When troubleshooting performance on a WLAN, an important statistic is *channel utilization*, as shown in Figure 40. Remember that RF is a shared medium and that 802.11 radios must take turns transmitting on any Wi-Fi channel. If the channel is oversaturated with 802.11 transmissions, performance will be negatively impacted. When they are not transmitting, both AP and client 802.11 radios listen to a channel every 9 microseconds for 802.11 transmission activity as well as non-802.11 transmissions.

FIGURE 40 Channel utilization



Here are some good channel utilization thresholds to live by:

- Eighty percent channel utilization impacts all 802.11 data transmissions.
- Fifty percent channel utilization impacts video traffic.
- Twenty percent channel utilization impacts voice traffic.

Monitoring and troubleshooting channel utilization is important because of perceived performance of the Wi-Fi network from end users. A common support call is that a user calls and complains that the Wi-Fi is slow. If the channel utilization is indeed over 80 percent, the Wi-Fi is in fact slow. Improper WLAN design with improper channel planning very often leads to CCI, which will cause high channel utilization. Oversaturation of clients and

high bandwidth applications can consume too much airtime on a channel, which is why proper capacity planning is important. Too many broadcast SSIDs, low basic data rates on APs, and any abundance of legacy clients are all airtime consumption culprits that will affect channel utilization.

The QBSS information element found in 802.11 beacon and probe response frames sent by APs is a good indicator of channel utilization from the perspective of the AP radio. The information found in the QBSS information element, as seen in Figure 41, is often used by WLAN vendor monitoring solutions and other applications to visualize channel utilization in the form of graphs or charts. Large enterprise customers mostly rely on the monitoring/troubleshooting capabilities from the perspective of the radio within an access point. RF statistics gathered from incoming client transmissions can also be centrally monitored. The information gathered is from the perspective of the AP radio.

FIGURE 41 QBSS information element

▼	Stations: 6, Channel Utilization: 48%		
Element ID:	11		
Length:	5 bytes		
Station Count:	6		
Channel Utilization:	124 (48%)		
Available Admission Capacity:	0		

In reality, the best view of the RF network will always be from the client perspective, which is why WLAN design and survey validation is so important. One approach that some enterprise WLAN vendors have taken is to have sensor APs act in place of client devices, whereby they log into other APs as client devices and then perform health checks. Keep in mind that clients have different receive sensitivities, and radios in APs are usually much more sensitive. Centralized monitoring and diagnostics using an AP radio is usually a great start; however, additional information may need to be collected from the client perspective when troubleshooting client issues.

Slow performance and bandwidth bottlenecks can indeed be a result of bad Wi-Fi design, and poor channel utilization is a good indicator of a Wi-Fi performance problem. However, the reason the Wi-Fi seems slow to the end user most often has nothing to do with the WLAN or channel utilization. Bandwidth bottlenecks are very often on the wired network due to poor wired network design. The number one bandwidth bottleneck is usually the WAN uplink at any remote site. But remember, the Wi-Fi will always be blamed first despite the inadequate WAN bandwidth.

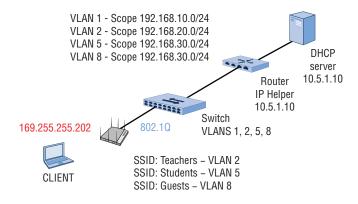
Layers 3–7 Troubleshooting

Although this booklet has focused on troubleshooting the WLAN at layers 1 and 2 of the OSI model, upper-layer trouble-shooting may still be a necessity. WLANs very often get blamed for causing problems that actually exist in the wired network at higher layers. If an employee cannot connect to the corporate WLAN, the employee will blame the WLAN even though the actual problem is somewhere else on the corporate network. If it can be determined that the problem is not a layer 1 or layer 2 problem, then the problem is usually a networking issue or problems with an application.

The good news is that many WLAN vendors offer upper-layer trouble-shooting tools that are available in network management servers or from the command line of APs. A common support call is that a user calls and complains that they have a Wi-Fi connection but cannot connect to the network. If you have already determined that the problem is not a Wi-Fi problem, move up the OSI stack to layer 3 to check for IP connectivity.

Please look at the diagram of a school WLAN in Figure 42; an AP is deployed in a school and is transmitting three SSIDS, one each for teachers, students, and guests. The teacher SSID is mapped to VLAN 2, the student SSID is mapped to VLAN 5, and the guest SSID is mapped to VLAN 8. The management interface of the AP is mapped to VLAN 1. All four VLANs are tagged across an 802.1Q trunk between the AP and the access switch. All four VLANs are mapped to respective subnets, and all IP addresses are supplied from defined scopes on the network DHCP server.

FIGURE 42 School WLAN diagram



As previously mentioned a common support call is that a user calls and complains that they have a Wi-Fi connection but cannot connect to the Internet. In this scenario, a student should be getting an IP address on the 192.168.30.0/24 network. A quick check determines that the student is connected the proper SSID but receives an *automatic private IP address* (APIPA) in the 169.254.0.0 to 169.254.255.255 range. This would be your first indication that the problem is most likely a wired-side network problem.

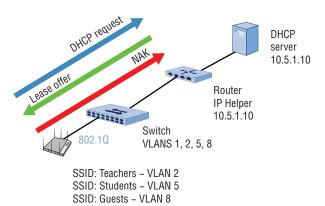
WLAN vendors might offer a diagnostic tool that can be used to report back if the VLANs are operational on the wired network as well as the subnet of each VLAN. As shown in Figure 43, an administrator can select an AP to perform a probe across a designated range of VLANs. Please note that VLAN 5 (the student VLAN) failed.

VI AN Probe Hostname: Coleman-AP250 The VLAN Probe is running! VLAN Range * to A (1-4094)Probe Retries (1-10)Timeout (1-60 seconds) Yas 192.168.10.0/24 2 192 168 20 0/24 Yes 3 No No No 6 No No 8 Yes 192.168.80.0/24

FIGURE 43 VLAN probe

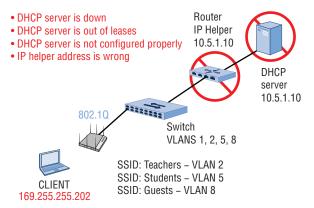
The diagnostic tool leverages the ability of the management interface of any access point to send out DHCP requests as shown in Figure 44. Once the probe starts, the management interface of the AP will send out multiple DHCP requests across all the designated VLANs. Each DHCP request is sent up the 802.1Q trunk and onto the wired network. Once the DHCP request finally reaches the DHCP server, a lease offer is sent back to the AP. The management interface of the AP does not need another IP address, therefore a NAK is sent back to the DHCP server. If the DHCP lease offer reaches the AP, then there is not an issue with the wired network. But if the DHCP lease offer does not reach the AP, then there is absolutely a wired-side problem and the diagnostic probe will show a negative result.

FIGURE 44 DHCP probe



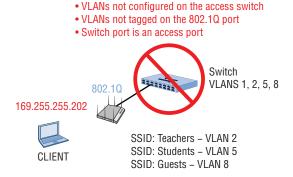
As shown in Figure 45, two common points of failure are the upstream router and the DHCP server. DHCP requests use a broadcast address and therefore an IP Helper (DHCP-Relay) address needs to be configured on the upstream router to convert the DHCP request into a unicast packet. If the router does not have the correct IP Helper address, then the DHCP request never makes it to the DHCP server. The DHCP server is more likely to be a point of failure. The DHCP server may have crashed, the scopes might not be configured correctly, or the server could simply be out of leases.

FIGURE 45 Backend DHCP failures



Although these two points of failure are certainly possible, the most likely culprit is the access switch, as shown in Figure 46. Almost 90 percent of the time the problem is an improperly configured access switch. The VLANs might not be configured on the switch, the VLANs might not be tagged on the 802.1Q trunk port, or the port has been misconfigured as an access port.

FIGURE 46 Misconfigured switch



Even if the WLAN clients are successfully receiving IP addresses, there could still be layer 3 network issues. Ping and traceroute/tracert commands are you next step to diagnosing your network. Ping and other available network query commands are available from every client operating system (OS) as well as the OS that runs on APs, switches, and routers.

Once you have determined that there is not a layer 3 networking problem, you can begin to investigate layers 4 through 7. Scott Adams created a funny *Dilbert* comic strip in 2013 that blames the firewall for all network problems: http://dilbert.com/strip/2013-04-07. This funny cartoon actually mirrors real life because an incorrectly configured firewall policy can be blocking TCP or UDP ports. In addition to stateful firewall capability, WLAN vendors have begun to build Application-layer firewalls capable of *deep packet inspection (DPI)* into access points. DPI provides visibility into applications being used over the WLAN, and Application-layer firewalls can block specific applications or groups of applications. Wherever the firewall may be deployed in your network, firewall log files may need to be reviewed if a higher layer problem is suspected.

Always remember that an access point is a wireless portal to complete network infrastructure. If the Wi-Fi network is not the problem, trouble-shooting layers 3–7 will be necessary.

WLAN Troubleshooting Tools

Although WLAN vendors provide significant diagnostic capabilities from their network management systems, every WLAN professional usually carries a wide array of toys in their personal WLAN troubleshooting toolkit. The following sections will describe some of the tools that are available.

WLAN Discovery Applications

To start troubleshooting WLANs, you will need an 802.11 client NIC and a WLAN discovery application such as WiFi Explorer, shown in Figure 47. WLAN discovery applications are a quick and easy way to give you a broad overview of an existing WLAN. WLAN discovery tools find existing Wi-Fi networks by sending out null probe request frames and listening for the 802.11 probe response frames and beacon frames sent by APs. Although a WLAN discovery tool will not give you the deep analysis that a protocol analyzer may provide, a lot of useful information can be gathered. For example, a WLAN discovery tool could immediately tell you that 80 MHz channels have been enabled on APs and performance is negatively impacted. A good WLAN discovery tool can give you a quick view of a number of transmitting APs and their channels, channel sizes, and security capabilities. Other available information includes signal strength, SNR, channel utilization statistics, and much more.

Numerous freeware and commercial discovery tools exist, including inSSIDer for Windows, WiFi Explorer for the macOS, and WiFi Analyzer for Android. You can download inSSIDer Office from www.metageek.net, WiFi Explorer from www.adriangranados.com, and WiFi Analyzer from bit.ly/WiFIAnalyze.

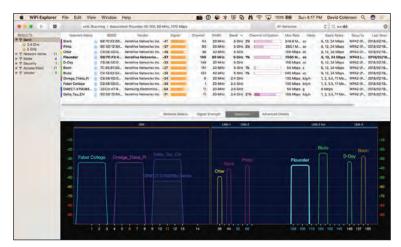
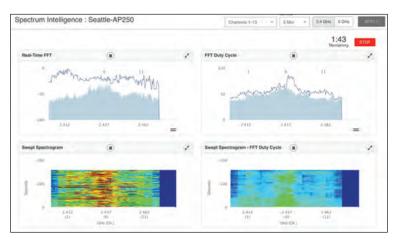


FIGURE 47 WLAN discovery tool

Spectrum Analyzers

Spectrum analyzers are frequency domain measurement devices that can measure the amplitude and frequency space of electromagnetic signals. Figure 48 depicts Aerohive's cloud-based spectrum analysis capabilities that can monitor in real time from the perspective of a remote Aerohive access point.





MetaGeek offers a USB-based adapter that is capable of monitoring both the 2.4 GHz and 5 GHz spectrums. MetaGeek's (www.metageek.net) Wi-Spy spectrum analyzer was used to identify the sources of RF interference shown earlier in this booklet in Figures 3 through 5. A spectrum analyzer is a layer 1 diagnostic tool that is most often used to identify sources of RF interference that originate from non-802.11 transmitters.

Protocol Analyzers

Protocol analyzers provide network visibility into exactly what traffic is traversing a network. Protocol analyzers capture and store network packets, providing you with a protocol decode for each packet captured, which is a readable display showing the individual fields and values for each packet. The power of a protocol analyzer is that it allows you to see conversations between various networking devices at many layers of the OSI model. Protocol analysis is sometimes the only way to troubleshoot a difficult problem. Many commercial WLAN protocol analyzers are available, such as Tamosoft's CommView for WiFi, www.tamos.com, as well as the popular freeware protocol analyzer Wireshark, www.wireshark.org.

Wired protocol analyzers are often called packet analyzers because they are used to troubleshoot IP packets that traverse wired networks. Remember, if the problem is not a layer 1 or layer 2 problem, the problem is not a Wi-Fi problem. Packet analysis of wired traffic is often necessary to troubleshoot problems that occur at layers 3 to 7.

WLAN protocol analysis is mostly used to look at layer 2 802.11 frame exchanges between APs and client devices. Wi-Fi radios communicate via 802.11 frame exchanges at the MAC sublayer. Unlike many wired network standards such as IEEE 802.3, which uses a single data frame type, the IEEE 802.11 standard defines three major frame types: management, control, and data. These frame types are further subdivided into multiple subtypes. When using a WLAN protocol analyzer to view 802.11 frame conversations, layers 3 to 7 will not be visible because encryption is enabled. Hopefully, all your 802.11 data traffic is encrypted.

A WLAN protocol analyzer and some WLAN discovery tools can also give you some insight to some layer 1 and RF statistical information. *Radiotap* headers provide additional link-layer information that is added to each 802.11 frame when they are captured. The drivers of an 802.11 radio supply additional information via the Radiotap header. Please understand that the Radiotap header is not part of the 802.11 frame format. However, the ability to see additional information, such as signal strength associated to each 802.11 frame heard by the WLAN protocol analyzer radio, is quite

useful. Wi-Fi expert Andrian Granados provides a more detailed explanation about the Radiotap header in his blog: www.adriangranados.com/blog/link-layer-header-types.

I always tell people that one of the best things I did early in my Wi-Fi career was to teach myself 802.11 frame analysis. Eighteen years later, 802.11 protocol analysis skills are still important and invaluable. Modernday WLAN protocol analyzer tools are more robust, but the 802.11 frame exchanges are constantly becoming more complex. As new 802.11 technologies such as 802.11ax become reality, the information inside 802.11 frame exchanges creates a Wi-Fi mosaic that is tough to interpret. No matter how much you might think you know about 802.11 frame analysis, you can always learn more.

Although there are many commercial WLAN protocol analyzers, Wireshark, www.wireshark.org, is open-source and the tool of choice for many WLAN professionals. The author of this booklet highly recommends two Wireshark video training courses created by Jerome Henry CWNE #45 and James Garringer CWNE #179. If you are a Wireshark novice, we highly recommend the *Wireshark Fundamentals LiveLessons* video training course (http://bit.ly/2BipIGs), which offers nearly five hours of instruction on using Wireshark to troubleshoot Ethernet and Wi-Fi networks and the protocols they transport. You will learn Wireshark capture basics, customization, filters, command-line options, and much more.

If you are interested in a deep-dive of 802.11 analysis, then we highly recommend the *Wireshark for Wireless LANs LiveLessons* video training course (http://bit.ly/2DX1swT), which offers more than eight hours of expert instruction on troubleshooting Wi-Fi networks using Wireshark. Nine lessons and sub-lessons will take you through learning the 802.11 MAC header, dissecting captured frames, advanced tools, and common WLAN problems that can be solved with proper analysis.

When using a protocol analyzer, using traffic filtering capabilities to focus on the networking conversations that you are trying to troubleshoot is essential. A reference guide for the most common 802.11 filters used in Wireshark is available for download at http://bit.ly/2mZzbyu courtesy of François Vergès, CWNE#180. Commercial protocol analyzers will have more advanced traffic filtering capabilities. Any good protocol analyzer will also have the ability to visualize traffic conversations as well as the ability to possibly provide intelligent diagnostics and suggested steps for remediation. Figure 49 is a screen capture from MetaGeek's EyePA WLAN protocol analyzer diagnosing an unacceptable percentage of layer 2 retransmissions on an AP along with suggestion steps to investigate the cause of the problem.

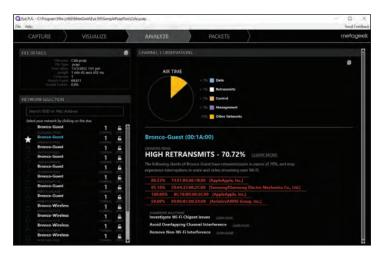


FIGURE 49 EyePA Analysis and Remediation

Identifying the correct location to place a network analyzer is an essential step in performing successful wireless network analysis. Incorrect placement of the WLAN protocol analyzer can lead to false conclusions being made. For example, if you are capturing traffic too far away from the source and destination, you might see a lot of corrupted frames; however, the intended recipient may not be experiencing any frame anomalies. An access point acts as the central point in an 802.11 wireless network, and all traffic must flow through the access point. Enterprise WLAN vendors offer direct packet capture from access points. In this scenario, if the analyzer reports a corrupted frame, it is more than likely that the AP also saw the frame as corrupted.

Throughput Test Tools

Throughput testing tools are used to evaluate bandwidth and performance throughout a network. Throughput testers normally work on a client/server model to measure data streams between two ends or in both directions. When testing downlink WLAN throughput, the 802.11 client should be configured as the server. When testing uplink WLAN throughput, the 802.11 client should be configured as the client communicating with a server behind the AP. *Iperf* is an open-source command-line utility that is commonly used to generate TCP or UDP data streams to test throughput. Many WLAN vendors offer Iperf as a CLI test utility from within the OS of access points. As shown in Figure 50, TamoSoft (www.tamos.com) offers a freeware GUI-based throughput tester that is available for Windows, macOS, iOS, and Android clients.

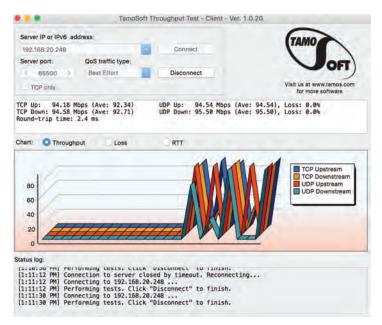


FIGURE 50 TamoSoft Throughput Tester

When performing throughput testing of the wireless link, always remember that you are not testing 802.11 data rates. Depending on WLAN network conditions the aggregate WLAN throughput is usually 50 percent of the advertised 802.11 data rate due to medium contention overhead. 802.11 data rates are not TCP throughput. The medium contention protocol of CSMA/CA consumes much of the available bandwidth. In laboratory conditions, the TCP throughput in an 802.11n/ac environment is 60 percent to 70 percent of data rate between one AP and one client. The aggregate throughput numbers are considerably less in real-world environments with active participation of multiple WLAN clients communicating through an AP.

Using client/server throughput test tools on the wired-side of the network is very often necessary. Remember, the reason the Wi-Fi seems slow to the end user most often has nothing to do with the WLAN or channel utilization. Bandwidth bottlenecks are very often on the wired network due to poor wired network design. Once again, the number one bandwidth bottleneck is usually the WAN uplink at any remote site.

Standard IP Network Test Commands

Always remember that you have standard network troubleshooting tools available within the various operating systems. Everyone knows that you always start

with Ping, the most commonly used network tool to provide a basic connectivity test between the requesting host and a destination host. Ping uses the *Internet Control Message Protocol (ICMP)* to send an echo packet to a destination host and listen for a response from the host. Use Ping to test IP connectivity between a WLAN client and a local network server. Use Ping to see if the WLAN client can reach the default gateway address. Ping the public Google DNS server at 8.8.8.8 to see if the WLAN client can access the Internet via the WAN. The following list includes some other commonly used network commands:

- Arp: The arp command is used to display the Address Resolution Protocol (ARP) cache, which is a mapping of IP addresses to MAC addresses. Every time a device's TCP/IP stack uses ARP to determine the MAC address for an IP address, it records the mapping in the ARP cache to speed up future ARP lookups. Viewing an ARP cache on access points is often helpful when troubleshooting.
- Tracert/Traceroute: The tracert or traceroute command is available is most operating systems to determine detailed information about the path to a destination host, including the route an IP packet takes, number of hops, and the response time between the various hops.
- Nslookup: The nslookup command is used to troubleshoot problems with *Domain Name System (DNS)* address resolution issues. DNS is used to resolve domain names to IP addresses. Use the nslookup command to look up a specific IP address associated with a domain name. Many WLAN captive web portals used for WLAN guest access rely on DNS redirection. If a WLAN captive web portal suddenly stops working, you should probably suspect that there is a DNS issue.
- Netstat: The netstat command displays network statistics for active TCP sessions for both incoming and outgoing ports, Ethernet statistics, IPv4 and IPv6 statistics, and more. Netstat is often useful when troubleshooting suspected application problems and firewall issues.

When troubleshooting from a WLAN client perspective, these commands will be easily available from the command line of devices running Windows, macOS, or Linux. Many freeware applications are also available for both iOS and Android so you can access these troubleshooting capabilities for smartphone and tablet mobile devices.

Are There Any CLI Commands to Troubleshoot WLAN Client Radios?

The simple answer is yes, depending on the operating system of the WLAN client device. The network shell (netsh) command can be used to configure and troubleshoot both wired and wireless network adapters

on a Windows computer. The netsh wlan show commands will expose detailed information about the Wi-Fi radio being used by the Windows computer. For example, netsh wlan show networks will display all the visible Wi-Fi networks that the client radio can see. A comparable command-line utility to configure and troubleshoot 802.11 network adapters on macOS computers is the airport command-line tool. Take the time to familiarize yourself with both the netsh wlan commands for Windows and airport commands for the macOS. As with any CLI command, execute the? command to view full options.

Secure Shell

When connecting to network hardware such as an access point or a switch, an SSH or serial client is required. *Secure Shell (SSH)* is used as the secure alternative to Telnet. SSH implements authentication and encryption using public-key cryptography of all network traffic traversing between the host and user device. The standard TCP port 22 has been assigned for the SSH protocol. Most WLAN infrastructure devices now support the second version of the SSH protocol, called SSH2. As a matter of policy, when WLAN devices are managed via the CLI, an SSH2-capable terminal emulation program should be used. Figure 51 shows the configuration screen of the popular freeware program PuTTY, which supports SSH2 and terminal emulation. PuTTY is often the freeware program of choice when having to climb a ladder and connect to the console port of an access point. Additionally, some operating systems such as macOS support SSH natively from the command line.

R PuTTY Configuration Category ∃ Session Basic options for your PuTTY session Logging Specify the destination you want to connect to □ Terminal Host Name (or IP address) Keyboard 10.1.1.10 22 Bell Features Appearance Load, save or delete a stored session Behaviour Saved Sessions Translation Selection Colours Default Settings □ Connection Data Save Proxy Delete Telnet Rlogin ⊕ SSH Serial Close window on exit: O Always O Never Only on clean exit Open. Cancel About.

FIGURE 51 PuTTY freeware SSH and serial client

Aerohive Troubleshooting Tools

Aerohive Networks's cloud-based network management solution, HiveManager, provides network administrators with the ability to easily manage, monitor, and troubleshoot WLANs. A huge advantage in cloud management is the ability to troubleshoot APs at remote locations. In this section, we will discuss some of the most commonly used HiveManager GUI-based diagnostic utilities as well as some of the commonly used CLI troubleshooting commands.

Diagnostic Utilities An administrator can select any monitored access point in HiveManager and run device diagnostic utilities shown in Figure 52. The remote Ping utility gives an admin the ability to send a ping from the management interface of an AP to the AP's default gateway or any other IP address. This gives you the ability to verify IP connectively and establish latency baselines on the wired side of the network. An administrator can also remotely view any AP's log and verify firmware versions and running configurations. The Show Roaming Cache diagnostic utility is helpful when troubleshooting roaming problems. An administrator also has the ability to verify and troubleshoot layer 3 roaming with the Show GRE Tunnel utility.

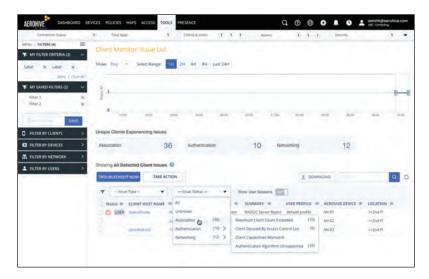
W UTILITIES Show Log Locked Clients Show Version RADIUS Test Show Running Config SSH Avallability Show Startup Config Device Information Show IP Routes Layer 2 Neighbor Information Show MAC Routes Locate Device Show ARP Cache Get Tech Data **Show Roaming Cache** VLAN Probe Show DNXP Neighbors **Device Diagnostics** Show DNXP Cache Device Client Information Unblind MAC/PPSK Bindings Show AMRP Tunnel Show GRE Tunnel REMOTE PACKET CAPTURE Show IKE Event

FIGURE 52 HiveManager device diagnostic utilities

Client Monitor Troubleshooting client connectivity to an access point is imperative. Client Monitor is a client connectivity tool that can proactively identify many of the PSK and 802.1X authentication problems discussed in

detail in this booklet. Client Monitor also provides descriptions of the identified problems as well as suggested remedies. As shown, in Figure 53, Client Monitor has an adjustable timeline that allows an administrator to focus on client connectivity issues as far back as 30 days. If you know a client MAC address, Client Monitor can also troubleshoot connectivity issues for active sessions between an AP and a WLAN client.





VLAN Probe The HiveManager VLAN Probe tool is the diagnostic tool that can often prove that the problem is not a Wi-Fi problem but instead a networking issue that exists on the wired network. As depicted earlier is this booklet in Figure 43, VLAN Probe can be used to report back if the VLANs are operational on the wired network as well as the subnet of each VLAN. The diagnostic tool leverages the ability of the management interface of any access point to send out DHCP requests as shown earlier in Figure 44. As seen in Figure 45 and Figure 46, VLAN Probe can help an administrator determine the actual points of failure on the wired network such as a DHCP server or improperly configured access switch.

RADIUS Test As you learned earlier, if an AP and a RADIUS server cannot communicate with each other, the entire 802.1X/EAP process will fail. As shown in Figure 26, the HiveManager RADIUS Test diagnostic utility gives the administrator the ability to test backend 802.1X/EAP communications between an AP and a RADIUS server.

Spectrum Intelligence As shown in Figure 48, the Spectrum Intelligence capability in Aerohive APs provides a live view of the RF environment so that you can plan for further WLAN deployment or troubleshoot WLAN issues such as high retransmission rates caused by RF interference. There are two main spectrum intelligence functions: providing a graphical rendering of the RF environment in a *fast Fourier transform (FFT)* trace and swept spectrogram, as well as identifying interfering devices such as cordless phones and microwave ovens.

Remote Packet Capture Aerohive partners with CloudShark, www.cloudshark.com, which is a web-based packet capture repository and suite of analysis tools that all work right in your browser. The Remote Packet Capture tool in HiveManager can capture 802.11 frames directly from the radios in any Aerohive AP. When an AP packet capture session is complete, the AP sends the capture results to HiveManager, which relays them to CloudShark for storage and display.

SSH Availability HiveManager provides a way to access devices remotely using the SSH protocol. Because best practices suggest that SSH access to internal devices be blocked from external sources, HiveManager does this by using an SSH proxy server to mediate the end-to-end connection between the AP to which you want to connect and your computer. This allows you to use your favorite terminal emulation program such as PuTTy to perform advance CLI troubleshooting commands.

The following access point CLI troubleshooting commands are among the most common:

- show capwap client: Aerohive APs communicate with HiveManager using the CAPWAP management protocol. This command allows you to observe the device discovery and connection states between an AP and HiveManager.
- show station: Allows you to quickly view all associated clients to an AP. You can see the client MAC addresses, IP addresses, connected upstream and downstream data rates, SNR, and more statistics.
- show interface: Displays all of the network interfaces of an AP, including the management interface, Ethernet interfaces, and radio interfaces. You can then drill down on any of the individual interfaces to view the detailed network interface information.
- show acsp neighbor: Displays RF information about all APs on every channel within listening range of the AP you are troubleshooting.
- show version detail: Displays the version of the HiveOS code that is running on the AP boot partition, as well as which version of the Aerohive operating system that resides on the AP backup partition.

Summary

Troubleshooting WLANs can be very challenging. Much of WLAN troubleshooting revolves around performance or connectivity issues that are a result of improper WLAN design. However, because of the ever-changing RF environment, problems such as roaming, hidden nodes, and interference are bound to surface. Never forget that Wi-Fi operates at layer 1 and layer 2 of the OSI model, and understand that the WLAN will always receive the blame no matter where the problem exists. Always remember to use troubleshooting best practices, analyze the problems at the different layers of the OSI model, and utilize all diagnostic tools that might be available.

WILEY END USER LICENSE AGREEMENT

Go to www.wiley.com/go/eula to access Wiley's ebook EULA.