

LEARNING MADE EASY

Aerohive Special Edition

Cloud-Managed Network Access Control (NAC)

for
dummies[®]
A Wiley Brand



Secure access for
all users and devices

Policy enforcement across
the enterprise network

Posture assessment
and remediation

Compliments
of



AEROHIVE[®]
NETWORKS

David Coleman – CWNE #4

About Aerohive Networks

Aerohive uses cloud management, machine learning, and artificial intelligence to radically simplify and secure the access network. Our cloud-managed wireless, switching, routing, and security technologies provide unrivalled flexibility in deployment, management, and licensing. Credited with pioneering controller-less Wi-Fi and cloud management, Aerohive delivers continuous innovation at cloud-speed that constantly challenges the industry norm, allowing customers to rethink what's possible. Our innovations and global cloud footprint radically simplify access network operation for 30,000+ customers and 10+ million daily users. See how at www.aerohive.com/customers.



Cloud-Managed Network Access Control (NAC)

Aerohive Special Edition

by David Coleman – CWNE #4

for
dummies[®]
A Wiley Brand

Cloud-Managed Network Access Control (NAC) For Dummies®, Aerohive Special Edition

Published by
John Wiley & Sons, Inc.
111 River St.
Hoboken, NJ 07030-5774
www.wiley.com

Copyright © 2019 by John Wiley & Sons, Inc., Hoboken, New Jersey

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 the prior written permission of the Publisher. 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>.

Trademarks: Wiley, For Dummies, the Dummies Man logo, The Dummies Way, Dummies.com, Making Everything Easier, and related trade dress 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. Aerohive and the Aerohive logo are trademarks or registered trademarks of Aerohive Networks. 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.

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 WEBSITE 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 WEBSITE MAY PROVIDE OR RECOMMENDATIONS IT MAY MAKE. FURTHER, READERS SHOULD BE AWARE THAT INTERNET WEBSITES 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, or how to create a custom *For Dummies* book for your business or organization, please contact our Business Development Department in the U.S. at 877-409-4177, contact info@dummies.biz, or visit www.wiley.com/go/custompub. For information about licensing the *For Dummies* brand for products or services, contact BrandedRights&Licenses@Wiley.com.

ISBN 978-1-119-56827-8 (pbk); ISBN 978-1-119-56826-1 (ebk)

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

Publisher's Acknowledgments

Some of the people who helped bring this book to market include the following:

Development Editor: Ryan Williams

Project Editor: Jennifer Bingham

Acquisitions Editor: Ashley Barth

Editorial Manager: Rev Mengle

Business Development

Representative: Karen Hattan

Production Editor: Siddique Shaik

Technical Editor: Faith Baynes

Marketing Review: Kathleen Hedde

Special Help from Aerohive:

Trevia Clark

Introduction

If you only had to keep one person away from a door, you would have an easy job (depending on the size of the door, but this is a hypothetical situation, so just roll with it). Now think about a multitude of people with a multitude of doors. It is a little more complicated. Apply that metaphor to your company's wired and wireless network, and it is even more complicated. And realistic.

Today, an ever-growing number of corporate and guest devices, as well as personal and IoT devices, need to be securely onboarded and provisioned with the access rights appropriate for their roles. Once authenticated and authorized, devices continue to pose security risks due to malware and compromised third-party applications. IT security departments need to be able to address these challenges cost efficiently, with a complete toolset for access management and control.

In this book, you learn about the key aspects of network access control (NAC) within enterprise IT networks. You also learn about Aerohive's solution: A³. Aerohive A³ is a *cloud-managed* NAC solution that secures all devices on the network — wireless and wired — including corporate, BYOD, guest, and IoT devices.

About This Book

Cloud-Managed Network Access Control (NAC) For Dummies, Aerohive Special Edition, consists of eight chapters:

- » Chapter 1 offers an overview of network access control technologies.
- » Chapter 2 shows the value of validating identity via authentication methods.
- » Chapter 3 analyzes onboarding security credentials onto client devices.
- » Chapter 4 looks at many aspects of providing network access for guests.
- » Chapter 5 reviews the importance of posture assessment.
- » Chapter 6 outlines methods of client device profiling.

- » Chapter 7 introduces you to Aerohive A³, a complete cloud-managed network access control (NAC) solution from Aerohive.
- » Chapter 8 lists ten key things to remember about cloud-managed NAC.

Foolish Assumptions

It has been said that most assumptions have outlived their usefulness, but I assume a few things nonetheless!

Mainly, I assume that you are an IT infrastructure professional — someone with networking, wireless, or security in your title — and that you work for a medium to large organization or enterprise with robust network access security requirements. Finally, you are interested in the many aspects and options that NAC has to offer.

If any of these assumptions describe you, then this book is for you! If none of these assumptions describe you, keep reading anyway. It is a great book and when you finish reading it, you will know a few things about the next generation of cloud-managed NAC!

Icons Used in This Book

Throughout this book, I occasionally use special icons to call attention to important information. Here is what to expect.



REMEMBER

This icon points out information you should commit to your non-volatile memory, your gray matter, or your noggin — along with anniversaries and birthdays!



TECHNICAL
STUFF

You will not find a map of the human genome here, but if you seek to attain the seventh level of NERD-vana, perk up! This icon explains the jargon beneath the jargon!



TIP

Tips are appreciated, never expected — and I sure hope you will appreciate these tips. This icon points out useful nuggets of information.



WARNING

These alerts point out the stuff your mother warned you about (well, probably not), but they do offer practical advice to help you avoid potentially costly or frustrating mistakes.

Beyond the Book

There is only so much I can cover in 80 short pages, so if you find yourself at the end of this book, thinking “Where can I learn more?” just go to www.aerohive.com. You can also read the latest edition of David Coleman and David Westcott’s books: *CWNA Certified Wireless Network Administrator Study Guide: Exam CWNA-107* (Sybex) and *CWSP Certified Wireless Security Professional Study Guide: Exam CWSP-205* (Sybex).

Where to Go from Here

If you do not know where you are going, any chapter will get you there — but Chapter 1 might be a good place to start! However, if you see a particular topic that piques your interest, feel free to jump ahead to that chapter. Each chapter is written to stand on its own, so you can read this book in any order that suits you (though I do not recommend upside-down).

IN THIS CHAPTER

- » Understanding the components in a network access control solution
- » Conducting posture assessments on your endpoint devices
- » Applying your policies to a multitude of devices

Chapter 1

Managing NAC from the Cloud

Because you are reading this book, you may or may not already be familiar with the concept of *network access control* (NAC). You can be proud if so, but it is time to catch everybody else up. The traditional definition of a NAC solution is a method of computer network security which combines endpoint integrity together with an access control solution such as 802.1X. However, NAC security has evolved to also encompass guest access management as well as BYOD and IoT device security. In this chapter, you take a look at the core components of a NAC solution, advanced NAC capabilities, and benefits of cloud-management.

Starting with the End

An *endpoint* is any computing device that requires access to a computer network for communications while connected. Examples of endpoints include smartphones, tablets, desktops, laptops, and IoT devices. In today's world, even a lightbulb can be an endpoint. Endpoint devices may require network access via either a wired or wireless connection.

A traditional NAC solution provides assessment checks of an endpoint to ensure that the devices have the required operating system, virus signatures, software patches, and other consideration. Additionally, access to an enterprise network is controlled using predefined role-based access policies together with an *authentication, authorization, and accounting* (AAA) solution. However, because the ways in which users and devices require network access has evolved over the years, NAC solutions have also had to evolve. Many NAC solutions now also offer enhanced access capabilities to include guest access management, *bring your own device* (BYOD) management, and security for IoT devices. NAC solutions have also grown to include device fingerprinting capabilities to correctly identify the make and OS of endpoint devices that are seeking connectivity.

With the growing acceptance of cloud computing, next-generation NAC solutions will also provide the ease and flexibility of cloud management and monitoring. In this chapter, I provide an overview of some of the basic concepts and components of NAC security. From here I will also point to other chapters for deeper dives into NAC security.

NAC Components

Even though NAC security has evolved, the primary goals of NAC remain the same. This section goes over what you can expect as part of your NAC solution.

Authentication, authorization, and accounting

Authentication authorization and accounting (AAA) is a key computer security concept that defines the protection of network resources.

Authentication

Authentication is the verification of identity and credentials. Users or devices must identify themselves and present credentials, such as usernames and passwords or digital certificates. More secure authentication systems use multifactor authentication, which requires at least two sets of different types of credentials to be presented.

Authorization

Authorization determines if the device or user is authorized to have access to network resources. This can include identifying whether you can have access based upon the type of device you are using (laptop, tablet, or smartphone), time of day restrictions, or location. Before authorization can be determined, proper authentication must occur.

Accounting

Accounting is tracking the use of network resources by users and devices. It is an important aspect of network security, used to keep a historical trail of who used what resource, when, and where. A record is kept of user identity, which resource was accessed, and at what time. Think of this process like signing in with a doorman. Keeping an accounting trail is often a requirement of many industry and government regulations, such as the *payment card industry* (PCI) security standard.



TECHNICAL
STUFF

The most common implementation of AAA security utilizes an 802.1X framework. All NAC solutions offer 802.1X security including support for RADIUS and a *public key infrastructure* (PKI). To learn more about 802.1X, please read Chapter 2.

Role-based access control and policy enforcement

Role-based access control (RBAC) is an approach to restricting system access to authorized users. After successful authentication, endpoint devices can be assigned defining network policies based on user roles, type of device, applications, time of day, location of the network, and other criteria. This process ensures that users can access only the resources they need and not other, potentially confidential and sensitive information. RADIUS attributes are often leveraged to assign different groups of users to different user traffic settings, including VLANs, firewall policies, bandwidth policies, and much more. Additionally, based on a change in the endpoint's network behavior or status, network access can be dynamically changed using RADIUS *Change of Authorization* (CoA). To assist in policy enforcement, most NAC solutions fully integrate with switches, firewalls, *mobile device management* (MDM)

solutions, *intrusion detection systems* (IDS), and third-party endpoint security applications.



TIP

To learn more about role-based access control and policy enforcement, please read Chapters 2 and 5.

Posture assessment

Posture assessment provides a set of rules to check the health and configuration of an endpoint and determine if it should be allowed access to the network. Posture assessment can be used to ensure endpoint integrity by validating up-to-date versions of device OS, antivirus, antimalware signatures, and application patches. Non-compliant endpoint devices can then be quarantined until they are updated. NAC solutions can often also assist the user during the remediation process necessary to bring the device into compliance. Noncompliant devices can be quarantined and if desired, IT can allow self-remediation of issues, by downloading the latest version of antivirus software. Self-remediation provides essential benefits in that it helps reduce IT workload.

Most posture assessments check the integrity on endpoint using preadmission checks before the endpoint can connect to the network. However, some NAC solutions also offer the capability to perform periodic post-admission checks after an endpoint has already joined the network. To learn more about posture assessment, please proceed to Chapter 5.

Advanced NAC Capabilities

Once you address the basic components of a NAC solution, you can expand into other capabilities.

BYOD

For many years, the primary purpose of enterprise WLANs was to provide wireless access for company-owned laptop computers used by employees. Some vertical markets, such as healthcare, retail, and manufacturing, also required WLAN access for company-owned mobile devices, such as VoWiFi phones and wireless barcode scanners. Over the last 10 years, however, there has been a massive

population explosion of Wi-Fi-enabled personal mobile devices. Wi-Fi radios are now the primary communications component in smartphones, tablets, PCs, and many other mobile devices.

Personal mobile Wi-Fi devices, such as smartphones and tablets, have been around for quite a few years. The Apple iPhone was first introduced in June 2007, and the first iPad debuted in April 2010. HTC introduced the first Android smartphone in October 2008.

These devices were originally meant for personal use, but in a very short time, employees wanted to also use their personal devices on company WLANs. Additionally, software developers began to create enterprise mobile business applications for smartphones and tablets. Businesses began to purchase and deploy tablets and smartphones to take advantage of these mobile enterprise applications. Tablets and smartphones provided the true mobility that employees and businesses desired, and within a few years, the number of mobile devices connecting to corporate WLANs surpassed the number of laptop connections. This trend is continuing, with many, if not most, devices shipping with Wi-Fi as the primary network adapter. Many laptop computers now ship without an Ethernet adapter because the laptop Wi-Fi radio is used for network access.

Although mobile devices initially were intended for personal use, organizations found ways of deploying corporate mobile devices with custom software to improve productivity or functionality. Employees also increasingly want to use their personal mobile devices in the workplace. Employees expect to be able to connect to a corporate WLAN with multiple personal mobile devices. The catchphrase *bring your own device* (BYOD) refers to the policy of permitting employees to bring personally owned mobile devices, such as smartphones, tablets, and laptops, to their workplace. A BYOD policy dictates which corporate resources can or cannot be accessed when employees connect to the company WLAN with their personal devices. Because of the proliferation of personal mobile devices, a BYOD policy is needed to define how employees' personal devices may access the corporate WLAN. A modern-day NAC solution must also be able to provide access based on predefined policies for BYOD endpoints in addition to company-owned devices.

Certificate onboarding

Many NAC solutions now also provide a means of certificate onboarding. Because 802.1X security requires X.509 digital certificates, a self-service device onboarding solution is often needed to automate the distribution and installation of certificates onto endpoint devices. The onboarding solution is most often used to install the root CA certificates on mobile devices to be used with an 802.1X/EAP-enabled SSID. Client certificates can also be provisioned with an onboarding solution. The main purpose of a self-service device onboarding solution is to provide an inexpensive and simple way to provision devices onto a secure corporate SSID. A device onboarding solution provides a self-service method for an employee to configure an endpoint and install security credentials. A NAC onboarding solution is often used both for company-owned endpoints as well as employee BYOD endpoints. To learn more, please proceed to Chapter 3.

Guest access

In today's world, business customers have come to expect guest WLAN access. Guest access can be a value-added service and often breeds customer loyalty. The primary purpose of a guest WLAN is to provide a wireless gateway to the Internet for company visitors and customers. Because guest users do not need access to company network resources, a guest WLAN must protect the company network infrastructure from the guest users. Think of these policies as providing a guest bedroom or bathroom so that nobody has to go through your stuff. Most NAC solutions now also offer advanced guest management features, including self-registration, social login, and captive web portals. To learn more about guest access, please read Chapter 4.

The Internet of Things

Over the years, most of the data generated on the Internet has been created by human beings. The theory of *Internet of Things* (IoT) is that in the future, the bulk of the data generated on the Internet might be created by sensors, monitors, and machines. It should be noted that 802.11 radio NICs used as client devices have begun to show up in many types of machines and solutions. Wi-Fi radios already exist in gaming devices, stereo systems, and video cameras. Appliance manufacturers are putting Wi-Fi NICs in washing machines, refrigerators, and automobiles. The use of

Wi-Fi radios in sensor and monitoring devices, as well as RFID, has many applications in numerous enterprise vertical markets. For example, wireless IoT sensors are used to control commercial lighting and HVAC systems.

Technology research firm 650 Group estimates that the number of wireless IoT endpoints in the enterprise market to grow by 80 percent over the next four years, reaching 8.6 billion units by 2023. Could this be the beginning of the self-aware Skynet predicted by the *Terminator* movies? All kidding aside, a large portion of IoT devices will most likely connect to the Internet with a Wi-Fi radio. Many IoT devices may also have an Ethernet networking interface in addition to a Wi-Fi interface. IT administrators must manage the onboarding, access, and security policies of IoT devices connecting to the corporate network. A good NAC solution with *device fingerprinting* capabilities is necessary to identify and classify IoT devices so that proper access policies can be enforced. To learn more about device fingerprinting, please read Chapter 6.

Cloud

Cloud computing and *cloud networking* are catchphrases used to describe the advantages of computer networking functionality when provided under a *Software as a Service* (SaaS) model. The *cloud* terminology essentially means a scalable private enterprise network that resides on the Internet. The idea behind cloud networking is that applications and network management, monitoring, functionality, and control are provided as a software service.

Amazon is a good example of a company that provides an elastic cloud-based IT infrastructure so that other companies can offer pay-as-you-go subscription pricing for enterprise applications and network services. Distributed computing models billed as “cloud” have progressed over time from a basic remote server to a truly scalable cloud-oriented architecture. *Elastic cloud* describes a cloud offering that provides variable service levels based on changing needs. A true elastic cloud is a distributed system where functions are provided by multiple microservices working collaboratively. Each microservice runs in its own processes and communicates through *application programming interfaces* (APIs). Services are hosted on pools of servers grouped by clusters in which any member can provide the same set of services as any other. The system scales by adding servers to clusters and removing them as needed. It achieves high performance and availability

by distributing loads among cluster members through the use of messaging queues and load balancers. True elastic cloud emphasizes openness and uses APIs extensively.

Next-generation NAC solutions such as Aerohive's A³ move NAC to the next-level with cloud management. A cloud-managed NAC solution scales globally by providing a central point of management for monitoring NAC enforcement nodes at multiple locations. A cloud-managed NAC can provide greater flexibility in terms of configuration for all sites and simplify workflows. In Chapter 7, I introduce you to Aerohive's A³ cloud-managed NAC solution.

IN THIS CHAPTER

- » Learning about the 802.1X framework
- » Finding out more about the RADIUS server
- » Leveraging attribute-value pairs

Chapter 2

Authentication Nation

At its heart, authentication just means making sure a device is actually what it says it is and not a sneaky bad device. A little simplistic, sure, but it is still the truth. A more realistic scenario demands a deeper dive into the alphabet soup of security protocols and standards. This chapter looks at the established standards that let devices connect to your network in the most secure and manageable method possible.

802.1X

NAC solutions are built on top of an 802.1X framework. The 802.1X standard is a *port-based access control* standard. The 802.1X-2001 standard was originally developed for 802.3 Ethernet networks. Later, 802.1X-2004 provided additional support for 802.11 wireless networks and Fiber Distributed Data Interface (FDDI) networks. The current version of the port-based access control standard, 802.1X-2010, defined further enhancements. 802.1X provides an authorization framework that allows or disallows traffic to pass through a port and thereby access network resources. 802.1X is a necessary component for authentication for wireless networking, although you can also use it in a wired environment.

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. A Layer 2 authentication protocol called *Extensible Authentication Protocol* (EAP) is used within the 802.1X framework to validate users at Layer 2. Figure 2-1 shows you these three components.




FIGURE 2-1: 802.1X framework.

Supplicant

A host with software that requests authentication and access to network resources is known as a *supplicant*. Each supplicant has unique authentication credentials that are verified by the authentication server. The supplicant uses an EAP protocol to present credentials to the authentication server to prove identity. Depending on which type of EAP protocol is used, the supplicant identity credentials can be in many different forms. However, usernames and their respective passwords are the most common form of identity information that is supplied by a supplicant to an authentication server.

In a WLAN, the supplicant is often the laptop or wireless handheld device trying to access the network. In a wired LAN, the supplicant is often a desktop computer.

Authenticator

An authenticator device blocks traffic or allows traffic to pass through its port entity. Authentication traffic is normally allowed to pass through the authenticator, while all other traffic is blocked until the identity of the supplicant has been verified. The authenticator maintains two virtual ports:



- » An *uncontrolled port* allows EAP authentication traffic to pass through.
- » A *controlled port* blocks all other traffic until the supplicant has been authenticated.

In a WLAN, the authenticator is usually either an AP or a WLAN controller. In a wired LAN, the authenticator is usually an access layer switch.

Authentication server

An authentication server (AS) validates the credentials of the supplicant that is requesting access and notifies the authenticator that the supplicant has been authorized. The authentication server will maintain a native database or may proxy query with an external database, such as an LDAP database, to authenticate the supplicant credentials. A RADIUS server normally functions as the authentication server.

What's the protocol, Kenneth?

Although the supplicant, authenticator, and authentication server work together to provide the framework for 802.1X port-based access control, an authentication protocol is needed to perform the authentication process. EAP is used to provide user authentication.

EAP is a flexible layer 2 authentication protocol used by the supplicant and the authentication server to communicate. The authenticator allows the layer 2 EAP traffic to pass through its virtual uncontrolled port. After the authentication server has verified the credentials of the supplicant, the server sends a message to the authenticator that the supplicant has been authenticated; the authenticator is then authorized to open the virtual controlled port and allow all other traffic to pass through. The entire authentication process occurs at layer 2. Once the controlled port is open, the supplicant has access to layers 3 through 7.

LDAP

Lightweight Directory Access Protocol (LDAP) is an application protocol for providing *directory services* over an IP network. The current version, LDAPv3, is defined in IETF RFC 4511. The directory service

infrastructure shares information about network resources such as files, folders, computers, users, groups, and so on. A directory service could be considered a database or data store, although it is different when compared to relational databases. Distributed directory information services use a hierarchical structure that can be accessed and managed using LDAP. In most IP networks, LDAP is used to provide access to a data store of usernames and passwords. Applications such as RADIUS (addressed in the “RADIUS” section later in this chapter) can be used to query an LDAP server to validate user or device credentials. LDAP sessions normally use TCP or UDP port 369. However, LDAP over SSL uses port 636.

Within an 802.1X authorization framework, the *authentication server* validates the credentials of a supplicant that is requesting access and notifies the authenticator that the supplicant has been authorized. The authentication server maintains a user database or may proxy with an external user database to authenticate user or device credentials. In almost all cases, a RADIUS server functions as the authentication server. The RADIUS server may hold a master native-user database but usually will instead query to a preexisting external LDAP database.

Any LDAP-compliant database can be queried by the RADIUS authentication server. Active Directory is the most commonly used external LDAP database, but a RADIUS server can also query LDAP-compliant databases such as eDirectory or OpenLDAP. Other flavors of LDAPv3-compliant directory services include Apple’s Open Directory and Apache Directory Server. Open Directory is the directory services framework used by macOS X and macOS X Server. Apache Directory Server is open source.

RADIUS

Remote Authentication Dial-in User Service (RADIUS) is a networking protocol that provides authentication, authorization, and accounting (AAA) capabilities for computers to connect to and use network services. RADIUS authentication and authorization is defined in IETF RFC 2865. Accounting is defined in IETF RFC 2866.

RADIUS was developed back in 1991 as a client/server authentication and accounting protocol. This protocol grew to be widely used by ISPs for dial-up users and later logically extended to VPN dial-up users (thankfully, the world has moved on from dial-up since then). RADIUS developed critical mass and had the capability to extend, or rather *broker*, authentication to many different user databases. This includes LDAP, Active Directory, SQL databases, flat files, and native RADIUS users.



TIP

RADIUS servers are sometimes referred to as AAA servers.

The RADIUS protocol has been around since 1991 and predates Wi-Fi technology. However, RADIUS is a key component of 802.1X security that is widely used in enterprise WLAN and wired LAN deployments. RADIUS servers almost always function as the authentication server that validates the credentials of a supplicant. Additionally, RADIUS attributes can be used to transport EAP frames and for server-based role assignment for user traffic settings.

RADIUS can fully integrate with any type of LDAP database. Because both RADIUS and LDAP are mature technologies, multiple deployment models exist for both scalability and redundancy.

All the RADIUS communications are between the RADIUS server and the AP, which is functioning as the RADIUS client. RADIUS packets are used to encapsulate EAP frames during the authentication exchange. Think of the RADIUS protocol as a transport mechanism for EAP authentication conversations between the supplicant and the authentication server.



REMEMBER

The RADIUS communications occur strictly between the RADIUS client (authenticator) and the RADIUS server.

RADIUS clients

The term *RADIUS clients* often get confused with Wi-Fi clients (supplicants). Instead, RADIUS clients are the network devices that communicate directly with a RADIUS server using the RADIUS protocol. RADIUS clients are the authenticators within an 802.1X/EAP framework. To make things even more confusing, RADIUS

clients are also sometimes referred to as the network access server (NAS).




TIP

When discussing 802.1X/EAP, the terms *authenticator*, *NAS*, and *RADIUS client* are all synonymous.

Shared secret

Don't confuse a shared secret with information told to a spouse. A *shared secret* is used between the authenticator and the authentication server for the RADIUS protocol exchange. As shown in Figure 2-2, Layer 2 EAP protocol communications occur between the supplicant and the authentication server. The 802.11 frames use what is called EAP over LAN (EAPOL) encapsulation between the supplicant and the authenticator to carry the EAP data. On the wired side, the RADIUS protocol is used between the authenticator and the RADIUS server. The EAP data is encapsulated within a RADIUS packet. A shared secret exists between the authenticator and the server, so that they can validate each other with the RADIUS protocol.



Source: Aerohive

FIGURE 2-2: Shared secret.

Configuration Gotchas

When configuring authenticators (such as access points and switches) and RADIUS servers, you usually encounter two configuration problems:

- » Nonmatching shared secrets
- » Wrong UDP ports

Ensure these values are correct before attempting any supplicant authentication attempts. RADIUS uses UDP ports 1812 for RADIUS authentication and 1813 for RADIUS accounting. These ports were officially assigned by the Internet Assigned Number Authority (IANA). However, prior to IANA allocation of UDP ports 1812 and 1813, many RADIUS servers used the UDP ports of 1645 and 1646 (authentication and accounting, respectively) by default.

RADIUS proxy and realms

The RADIUS protocol can make use of networking *realms*, which identify where the RADIUS server should forward the RADIUS requests across the Internet and between ISPs. Realm naming formats are defined in RFC 4282. Realms defined by J.R.R. Tolkien and George R.R. Martin should not be honored here.

A RADIUS server can be a proxy to one or more centralized RADIUS servers. A username can use one of these two forms:

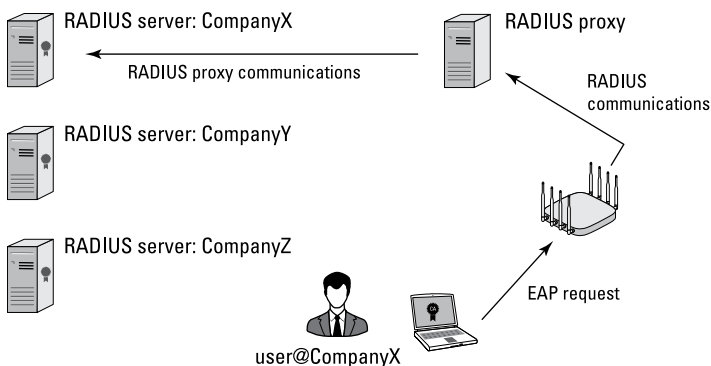
- » DOMAIN/username
- » username@domain

This valuable information can tell the first authentication server which final destination authentication server to request authentication from. In this case, *domain* is synonymous with the term realm as referred to previously.

For example, if an AP was configured in a regional office or subsidiary of a large enterprise, the AP might point to an authentication server located at that facility or perhaps in a nearby datacenter. Suppose an employee with a common name (for example, Kathleen) from the parent company was traveling on a business trip and visiting the subsidiary's office. When Kathleen logged in with Aerohive/kathleen, the authentication would know that it needed to contact the remote Aerohive user account RADIUS server. In this situation, the first authentication server that the AP pointed to would have just performed a proxy authentication to the final authentication server where Kathleen's user account resided.

When a domain is used, it is commonly referred to as a *realm-based authentication*. This method of authentication allows a user to authenticate to a realm or subrealm. When a RADIUS server receives an AAA request for a username containing a realm, the

server will reference a table of configured realms. If the realm is known, the server will proxy the request to the home RADIUS server for that domain. As shown in Figure 2-3, the realm in this case is the domain being sent in the supplicant identity. Authentication of the user is directed to a RADIUS server for CompanyX based on the domain value supplied. A global example of realm-based authentication is *eduroam* (education roaming).



Source: Aerohive

FIGURE 2-3: Realm-based authentication.

RadSec

The RADIUS protocol depends on the unreliable transport protocol UDP using port 1812 and has some potential security weaknesses. RADIUS is based on the MD5 algorithm, which has been proven to be insecure. Therefore, the IETF RFC 6614 defines Transport Layer Security (TLS) for RADIUS. The defined protocol is often called RadSec or *RADIUS over TLS*. RadSec is the next-generation RADIUS transport that relies on TCP and TLS for reliable and secure transport with integrity verification.

The main focus of RadSec is to provide a means of securing the communication between RADIUS/TCP peers using TLS encryption. The default destination port number for RadSec is TCP/2083. Authentication, accounting, and dynamic authorization changes do not require separate ports. RadSec can be used for RADIUS packets that traverse through different administrative domains and networks. The academia roaming access service, eduroam, has already begun to use RadSec globally.

EDUROAM

A global example of realm-based authentication is *eduroam* (education roaming). The technology behind eduroam is based on the IEEE 802.1X standard and a hierarchy of RADIUS proxy servers. Eduroam is a secure, worldwide roaming access service developed for the international research and higher education community. Eduroam allows students, researchers, and staff from participating institutions to obtain Internet connectivity across campus and when visiting other participating institutions. Authentication of users is performed by their home institution, using the same credentials as when they access the network locally. Depending on local policies at the visited institutions, eduroam participants may also have additional resources at their disposal. The eduroam Architecture for Network Roaming is defined in RFC 7593. More information about the eduroam service can be found at www.eduroam.org.

Attribute-value pairs

An attribute is a portion of information that determines the properties of a field or tag in a database. Attributes usually come in name and value pairs like `name="value"`. An *attribute-value pair* (AVP) is a representation of data in computer systems and applications. An attribute-value pair can be used to store and provide data in a database — for example, an attribute called `last-name` followed by its value pair, which is the actual last name of a person. The IETF designates an original set of 255 standard RADIUS attributes that can be used to communicate AAA information between a RADIUS client and a RADIUS server. The attribute-value pairs (AVPs) carry data in the RADIUS request and response packets. Figure 2-4 shows a packet capture of a RADIUS Access-Accept packet with a series of AVPs.

Vendor-specific attributes

RADIUS vendor-specific attributes (VSAs) are derived from the IETF attribute (26) `Vendor-Specific`. This attribute allows a vendor to create an additional 255 attributes however it wants. So many options!


```

RADIUS Protocol
Code: Access-Accept (2)
Packet identifier: 0x6 (6)
Length: 97
Authenticator: fbba6a784c7dec314caf0f27944a37b
[Time from request: 0.000114000 seconds]
Attribute Value Pairs
  AVP: l=6 t=Framed-IP-Address(8): Assigned
  AVP: l=6 t=Framed-MTU(12): 576
  AVP: l=6 t=Service-Type(6): Framed(2)
  AVP: l=21 t=Reply-Message(18): Hello
  AVP: l=6 t=EAP-Message(79) Last Segment[1]
    EAP fragment
      Extensible Authentication Protocol
        Code: Success (3)
        Id: 1
        Length: 4
  AVP: l=18 t=Message-Authenticator(80): b9c4ae6213a71d32125ef7ca4e
    Message-Authenticator: b9c4ae6213a71d32125ef7ca4e4c6360

```

Source: Aerohive

FIGURE 2-4: RADIUS attribute-value pairs.

Data that is not defined in standard IETF RADIUS attributes can be encapsulated in the (26)Vendor-Specific attribute. This standard allows vendors to support their own extended attributes otherwise not suitable for general use. VSAs allow RADIUS client vendors, such as the manufacturers of access points and switches, to support their own proprietary RADIUS attributes. The vendors typically offer free VSA dictionary files, which can be easily imported into most popular commercial RADIUS servers as well as the open source FreeRADIUS.

VLAN assignment

A common WLAN design strategy links a single user VLAN to a unique SSID. Most WLAN vendors allow a radio to broadcast as many as 16 SSIDs. However, broadcasting out 16 SSIDs is a bad practice because of the Layer 2 overhead created by the 802.11 management and control frames for each SSID. The broadcast of 16 SSIDs results in degraded performance due to MAC layer overhead. The best practice is to never broadcast out more than 3 or 4 SSIDs.

What if you need your employees segmented into different VLANs? Can a single employee SSID be mapped to multiple VLANs? When a RADIUS server queries Active Directory, the database responds with the memberOf LDAP attribute. The value of this attribute includes a list of any Active Directory Groups to which a user belongs. For example, this information can then be used in

RADIUS access policies. RADIUS servers can be configured with different access policies for different groups of users. The RADIUS access policies are usually mapped to different LDAP groups.

Properly configured access policies can then leverage RADIUS attributes for VLAN assignment when using 802.1X authentication on the employee SSID. When a RADIUS server provides a successful response to an authentication request from a supplicant, the Access-Accept response can contain a series of AVPs. One of the most popular uses of RADIUS AVPs assigns users to VLANs based on the identity of the authenticating user. Instead of segmenting users to different SSIDs that are each mapped to a unique user VLAN, all the users can be associated to a single SSID and assigned to different VLANs.



TIP

This same strategy can be used when dynamically assigning VLANs to access ports on a managed switch for wired clients. The VLAN ID or VLAN name must be configured on the switch acting as 802.1X authenticator, and must match the VLAN ID or VLAN name sent by the RADIUS server. Policies can also be created that move a wired device to a guest VLAN based on a failed authentication.

RADIUS server attributes used for dynamic VLAN assignment are defined in RFC 2868.

- » (64) **Tunnel-Type:** The value should be set to VLAN.
- » (65) **Tunnel-Medium-Type:** The value should be set to IEEE-802.
- » (81) **Tunnel-Private-Group-ID:** The value should be set to the VLAN ID or the VLAN name.

Role-based access control

Using RADIUS attributes for user VLAN assignment has been a network design strategy for many years. However, RADIUS attributes can be further leveraged to assign different groups of users to all kinds of different user traffic settings, including VLANs, firewall policies, bandwidth policies, and much more.

Role-based access control (RBAC) is an approach to restricting system access to authorized users. The RBAC approach includes three main components:

- » Users
- » Roles
- » Permissions

Separate roles can be created, such as the sales role or the marketing role. You can define user traffic permissions in these ways:

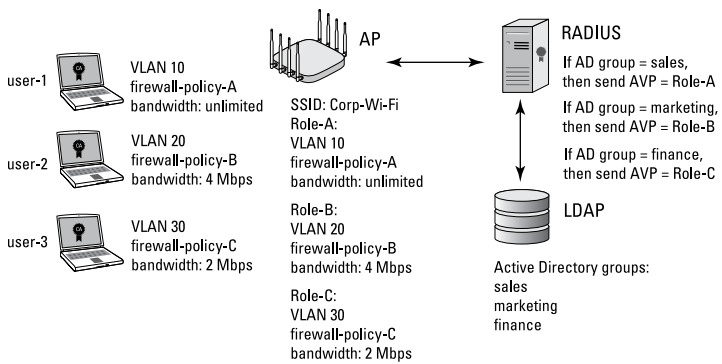
- » Layer 2 permissions (MAC filters)
- » VLANs
- » Layer 3 permissions (access control lists)
- » Layers 4 through 7 permissions (stateful firewall rules)
- » Bandwidth permissions

All these permissions can also be time-based. The user traffic permissions are mapped to the roles. Some vendors use the term *roles*, whereas other vendors, such as Aerohive Networks, use the term *user profiles*.

When a user authenticates using 802.1X, RADIUS attributes can be used to assign users to a specific role automatically. All users can associate to the same SSID but be assigned to unique roles. This method is often used to assign users from certain Active Directory groups into predefined roles created on a WLAN controller or access point. Each role has unique access restrictions. Once users are assigned to roles, they inherit the user traffic permissions of whatever roles they have been assigned.

Figure 2-5 depicts a RADIUS server with three unique access policies mapped to three different Active Directory (AD) groups.

For example, user-2 belongs to the marketing AD group. Based on the RADIUS access policy for that AD group, when user-2 authenticates, the RADIUS server will send the AP a RADIUS packet with an attribute that contains a value relevant to Role-B, which has been configured on the AP. The user-2 WLAN client will then be assigned to VLAN 20, firewall-policy-B, and a bandwidth policy of 4 Mbps.



Source: Aerohive

FIGURE 2-5: Assigning RADIUS roles with attributes.

Networking vendors may use VSAs for role assignment, but the standard IETF RADIUS attribute (11) `Filter-Id` is also often used for role assignment of user traffic permissions.

Machine Authentication

Although 802.1X/EAP is most often used to authenticate and authorize network access for users, computer devices can also be authorized. *Machine authentication*, also known as computer authentication, is the concept of ensuring the device requesting access to the network is authorized in a separate but chained authentication process. Machine authentication is often deployed as an extra layer of security in Active Directory environments.

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. The information stored in this SID is unique for each AD machine.

The computer account is used to identify the machine, even when no user is logged in, which can be used to provide the machine access to the network. Usually the machine does not need full access to the entire network and is often very restricted. Machine authentication is usually more about validating through 802.1X/EAP that the computer is authorized on the corporate network. 802.1X/EAP user authentication can then be used to grant further

access. A computer might be placed into a unique VLAN after machine authentication and is transitioned to a different VLAN after user authentication. The machine authenticates using the cached machine credentials and the user authenticates when logging into the domain. Machine authentication is often used in enterprise environments where the computer is shared by multiple employees with different user accounts.

Many RADIUS servers are designed to note the multiple accounts authenticating from the same MAC and use this to chain the authentication process together to provide multiple authentication factors that are used to evaluate the security level of the device. For example, a person with a user account logging into a network with their own device could be put into a more restricted VLAN than if they had logged in with a corporate-approved Windows computer.

Note that machine authentication with 802.1X/EAP and Active Directory was designed for devices using the Windows OS. When a user logs into a Windows system, the context of the system on the network changes and a new 802.1X/EAP authentication occurs. Your options for machine authentication with non-Windows operating systems are limited and complex. Macs have three different modes for configuring authentication:

- » System
- » Login Window
- » User

However, Macs cannot use multiple modes and switch context as in Windows. Machine authentication is also not really an option for mobile devices using iOS and Android OS.

Captive Web Portal and MAC Authentication

Most of this chapter revolved around using RADIUS with 802.1X authentication for strong security. However, a RADIUS server can also be used to authorize users and devices with weaker

authentication methods, such as *captive web portal* authentication and MAC address authentication.

RADIUS servers are often used with a captive portal web page to authenticate guest users for guest WLANs. A captive web portal solution will query a RADIUS server with a username and password using a weak authentication protocol such as MS-CHAPv2. The native database of the RADIUS server is normally used to validate the guest users. Captive web portal authentication is also often used together with *bring your own device* (BYOD) solutions for validating employee credentials. The employee database would most likely be Active Directory, which would in turn be queried by the RADIUS server.

- » Sorting out EAP protocols
- » Learning about digital certificates
- » Allowing self-service onboarding

Chapter 3

All Aboard for Onboarding

To keep everything secure, you need the right credentials. In this case, secure authentication using 802.1X requires digital certificates. But how do you easily distribute and properly install the correct certificates to all the supplicants? This chapter examines the role certificates serve within an 802.1X security framework and the means necessary to provision certificates onto client devices.

EAP Types

EAP stands for Extensible Authentication Protocol. The key word in EAP is *extensible*. EAP is a layer 2 protocol that is very flexible, and many different flavors of EAP exist. *Protected Extensible Authentication Protocol* (PEAP) is probably the most common and most widely supported EAP method used in enterprise 802.1X security. Most types of EAP require a server-side digital certificate to validate the authentication server. A server-side certificate is installed on the RADIUS server, while the *certificate authority* (CA) root certificate resides on the supplicant. During the EAP exchange, the supplicant's root certificate is used to verify the server-side

certificate. As depicted in Figure 3-1, the certificate exchange also creates an encrypted Secure Sockets Layer (SSL)/Transport Layer Security (TLS) tunnel in which the supplicant's username/password credentials or client certificate can be exchanged. Many of the secure forms of EAP use tunneled authentication. The SSL/TLS tunnel is used to encrypt and protect the user credentials during the EAP exchange.

Server certificates and root CA certificates

As shown in Figure 3-1, a server certificate must be created and installed on the authentication server.




FIGURE 3-1: Server certificate and root CA certificate.

Additionally, the root Certificate Authority (CA) public certificate that was used to create the server certificate must be installed on the supplicants. Distribution and installation of the root CA certificate to multiple WLAN supplicants is often the biggest challenge when deploying 802.1X/EAP security.

The authentication server certificate, in conjunction with the root CA public certificate, serves two major purposes.

Validating the authentication server

The server-side certificate is first used to validate the identity of the server to the supplicant. This process is akin to the supplicant saying, “Oh, I know who you are,” before the supplicant submits its own sensitive identity information (it is a little more technical, though).

During the EAP exchange, the RADIUS server sends the server certificate to the supplicant. The server certificate is then validated by the root CA certificate that resides on the supplicant. This is possible because the server certificate was created and signed by the root CA.


WHAT IS THE DIFFERENCE BETWEEN SSL AND TLS?

Essentially there is no difference between SSL and TLS. The original term Secure Sockets Layer (SSL) refers to an earlier version of the Transport Layer Security (TLS) protocol that is used to provide end-to-end encryption. TLS uses symmetric cryptography to provide bidirectional encryption between two computer applications. The keys for this symmetric encryption are generated uniquely for each connection and are based on a shared secret negotiated using a TLS handshake. The identity of the two devices using TLS can be authenticated using public-key cryptography, thus the need for certificates. The term SSL is still often used to refer to TLS. For example, HTTPS is HTTP-within-SSL/TLS. SSL (TLS) establishes an encrypted bidirectional tunnel for data between two hosts. HTTP is an application protocol that is the foundation of data communication on the World Wide Web. When HTTP is communicated via an SSL/TLS tunnel, it is known as HTTPS.

Creating an encrypted TLS tunnel

EAP protocols that require a server-side certificate for the authentication server are used to create TLS encryption tunnels. TLS is a cryptographic protocol normally used to provide secure communications at the Transport layer of the OSI model. However, in the case of 802.1X/EAP, TLS technology is leveraged at Layer 2. Similar to a browser-based HTTPS session, the TLS protocol uses end-to-end encryption. The exchange of the username/password credentials are protected within the SSL/TLS tunnel.

So where does the server certificate get created before it is installed on the RADIUS server? The simple method is to purchase a server certificate from a trusted root Certificate Authority (CA), such as GoDaddy (www.godaddy.com) or Verisign (www.verisign.com). The server certificates usually cost about several hundred dollars a year for each RADIUS server you deploy. The good news is that the root CA certificate already resides on most devices that can function as applicants, as seen in Figure 3-2.

| AAA Certificate Services | | | | |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|---------------------------|--------------|--|
|  Root certificate authority Expires: Sunday, December 31, 2028 at 3:59:59 PM Pacific Standard Time ● This certificate is valid | | | | |
| Name | Kind | Expires | Keychain | |
| Entrust Root Certification Authority - E-C1 | certificate | Dec 18, 2037, 7:55:36 AM | System Roots | |
| Entrust Root Certification Authority - G2 | certificate | Dec 7, 2030, 9:55:54 AM | System Roots | |
| Entrust.net Certification Authority (2048) | certificate | Dec 24, 2019, 10:20:51 AM | System Roots | |
| Entrust.net Certification Authority (2048) | certificate | Jul 24, 2029, 7:15:12 AM | System Roots | |
| ePKI Root Certification Authority | certificate | Dec 19, 2034, 6:31:27 PM | System Roots | |
| Federal Common Policy CA | certificate | Dec 1, 2030, 8:45:27 AM | System Roots | |
| GeoTrust Global CA | certificate | May 20, 2022, 9:00:00 PM | System Roots | |
| GeoTrust Primary Certification Authority | certificate | Jul 16, 2036, 4:59:59 PM | System Roots | |
| GeoTrust Primary Certification Authority - G2 | certificate | Jan 18, 2038, 3:59:59 PM | System Roots | |
| GeoTrust Primary Certification Authority - G3 | certificate | Dec 1, 2037, 3:59:59 PM | System Roots | |
| Global Chambersign Root | certificate | Sep 30, 2037, 9:14:18 AM | System Roots | |
| Global Chambersign Root - 2008 | certificate | Jul 31, 2038, 5:31:40 AM | System Roots | |
| GlobalSign | certificate | Mar 18, 2029, 3:00:00 AM | System Roots | |
| GlobalSign | certificate | Jan 18, 2038, 7:14:07 PM | System Roots | |
| GlobalSign | certificate | Jan 18, 2038, 7:14:07 PM | System Roots | |
| GlobalSign | certificate | Dec 15, 2021, 12:00:00 AM | System Roots | |
| GlobalSign Root CA | certificate | Jan 28, 2028, 4:00:00 AM | System Roots | |
| Go Daddy Class 2 Certification Authority | certificate | Jun 29, 2034, 10:06:20 AM | System Roots | |
| Go Daddy Root Certificate Authority - G2 | certificate | Dec 31, 2037, 3:59:59 PM | System Roots | |

Source: Aerohive

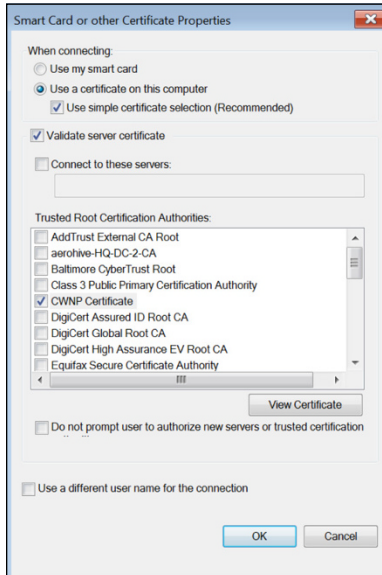
FIGURE 3-2: Trusted root CA certificates.

The major trusted CAs pay a lot of money to have their public root certificates accessible within the various operating systems. The main advantage of purchasing a server certificate from a trusted CA is that there is no need to distribute and install root certificates on wireless or wired clients because they already are there. It is good to know people.

One downside of using a public CA with 802.1X is that an attacker can possibly perform a man-in-the-middle attack. An attacker can use a rogue AP along with rogue RADIUS server and a server certificate that was also created from the same public CA. This attack is complex and has many moving parts. However, because the chain of trust might be compromised, many organizations choose to instead create a server certificate signed by an internal private CA, such as Microsoft Certificate Services. Much like a public CA, a private CA establishes an internal company trust chain using separate certificates for the root and the servers. Many companies choose this method because they prefer to keep all the security in-house and therefore install the server certificate signed by an internal CA on the RADIUS server. The main challenge with using an internal private CA is the distribution and installation of the root CA certificate from the internal CA to all the supplicants.

Client Certificates

A common mistake that people make is to confuse the root CA certificate that is installed on the supplicants with client-side certificates. A client certificate is an entirely different animal within a PKI infrastructure. As shown in Figure 3-3, client certificates can also be installed on a WLAN supplicant and be used as client credentials with some types of EAP authentication.



Source: Aerohive

FIGURE 3-3: The supplicant client certificate.

The most commonly deployed protocol that uses client certificates is *EAP-Transport Layer Security (EAP-TLS)*. Server certificates and a root CA certificate are still used in conjunction to validate the RADIUS server. However, the client certificate is used as the validation credential for the supplicant. Adding client certificates into the mix with 802.1X security provides an extra level of security but also adds an extra level of management and cost. Client certificates or smart cards with EAP-TLS security are often used

in military, government, and financial institutions. PEAP is more widely deployed for enterprise solutions, because the protocol does not require client certificates and is easier to manage. However, the use of EAP-TLS for 802.1X security is growing in the enterprise.

Using a public CA is usually cost-prohibitive because every client certificate costs several hundred dollars. Therefore, a private CA with an internal PKI is used to create and manage the client certificates. Management of client-side certificates requires much more time and proper skill sets, which many administrators might not have. Furthermore, the client certificates need to be provisioned onto the clients along with the root CA certificates. The provisioning of client certificates to company devices can be automated via a GPO in Windows. However, challenges also remain for distribution of client certificates for employee personal devices and devices using a non-Windows OS. Which is a LOT of devices.

Onboarding

One of the biggest challenges for an 802.1X/EAP solution is distributing and installing certificates onto client devices. Using an internal PKI infrastructure as a private CA first requires the creation of an original root certificate. Server certificates are then created and can be signed by the root CA. The recommended design for an internal PKI is implementing two or three tiers. At minimum, that scenario includes a root CA and an issuing CA with a possible intermediate CA in between those.

The server certificates are installed on the RADIUS servers. However, you still need to distribute and install the root certificate to all the wired and wireless supplicants. For example, the root certificate must be installed in the Trusted Root Certification Authorities Store of a Windows machine. Installing the root certificate onto Windows laptops can be easily automated using a Group Policy Object (GPO) if the Windows laptop is part of the Active Directory (AD) domain. However, a GPO cannot be used for macOS, iOS, or Android mobile devices, or for personal Windows BYOD devices that are not joined to the AD domain. Manually

installing certificates on mobile devices and employee-owned devices is a complex and time-consuming task.

For this reason, some companies deploy mobile device management (MDM) solutions. An MDM solution uses an encrypted over-the-air provisioning of certificates during the MDM enrollment process. One aspect of an MDM profile allows for provisioning Wi-Fi settings. You can lock down company-owned devices with a specific Wi-Fi profile that designates the corporate SSID and proper security settings. An MDM profile can also be used to deploy Wi-Fi settings to an employee's personal device. An MDM solution is an effective way to securely provision root CA certificates on mobile devices. Client certificates must also be provisioned if EAP-TLS is the chosen 802.1X security protocol. Some large enterprise companies use an MDM solution solely for the purpose of onboarding certificates to WLAN client devices because of the wide variance of operating systems. Many network access control (NAC) solutions such as Aerohive's A³, can integrate with third-party MDM systems.

Self-service device onboarding for employees

Instead of a full-blown and expensive MDM solution, another option is a *self-service device onboarding* solution for certificates. Self-service device onboarding solutions are typically much cheaper and simpler to deploy as an employee BYOD solution. Self-service onboarding solutions are used primarily to provision employee devices and are not used to enforce MDM device restrictions. MDM privacy concerns are no longer an issue for the employee personal devices when using a stand-alone certificate onboarding solution.

Vendors' self-service solutions, such as Aerohive's A³, enable employees to easily self-install security credentials, such as an 802.1X root CA certificate. EAP-TLS and client certificates add another layer of administrative complexity. Client certificates can also be provisioned via a self-service onboarding solution.



REMEMBER

The main purpose of a self-service device onboarding solution is to provide an inexpensive and simple way to provision employee personal WLAN devices onto a secure corporate SSID. A self-service device onboarding solution is not meant to offer all the

monitoring and restriction aspects of a full-blown MDM. Instead, a device onboarding solution provides a self-service method for an employee to configure the BYOD supplicant and install security credentials, such as an 802.1X root CA certificate.

Consider this scenario: Daniel logs in to the enterprise network as an employee from his corporate computer using 802.1X/EAP. His username is verified in the LDAP database using RADIUS. In this scenario, Daniel is trusted as the user because his username and password are valid. His corporate laptop is trusted because his machine can be validated. However, Daniel using his corporate laptop is different from Daniel using his smartphone, Daniel using his personal laptop, or Daniel using his tablet.

In the world of authentication and encryption, 802.1X/EAP is typically the required method to provide secure access to corporate networks and data. However, configuration of the client supplicant is typically not a task that can be easily performed by a non-technical user. Additionally, the root CA certificate needs to be securely transferred to the client device and installed. This can be a problem for corporations and employee BYOD users, since the corporation will not provide access without a properly configured device. As you can imagine, requiring a trained IT help desk person to configure employee personal devices is not practical. The solution is a process known as *onboarding*.

A properly configured and secured 802.1X/EAP network requires that a root CA certificate be installed on the supplicant. As mentioned earlier in this chapter, installing the root certificate onto Windows laptops can be easily automated using a Group Policy Object (GPO) if the Windows laptop is part of the Active Directory (AD) domain. However, a GPO cannot be used for macOS, iOS, or Android mobile devices, or for personal Windows BYOD devices that are not joined to the AD domain. Manually installing certificates on mobile devices and employee-owned devices is an administrative nightmare.

The onboarding solution is most often used to install the root CA certificates on mobile devices to be used with an 802.1X/EAP-enabled SSID. Client certificates can also be provisioned with an onboarding solution. Some of the Wi-Fi vendors that offer per-user PSK solutions also offer onboarding solutions that can

provision mobile devices with Wi-Fi client profiles configured with unique individual PSKs.

Self-service onboarding solutions for personal employee devices can come in many different forms and can be either vendor-specific or vendor-agnostic. Onboarding solutions may use an application that uses similar over-the-air provisioning aspects of MDM solutions to securely install certificates and Wi-Fi client profiles onto mobile devices. Self-service onboarding solutions may also use custom applications built using a WLAN vendor's *application programming interface* (API). Regardless of the solution, the device onboarding normally requires an initial WLAN connection to complete the self-service process.

Dual-SSID onboarding

Dual-SSID onboarding is performed using an open SSID and an 802.1X/EAP secure corporate SSID. The employee initially connects to the open SSID and will be prompted with a captive portal page. Depending on how onboarding is implemented, the employee may log in directly through the captive portal using their corporate username and password, or the employee may click a link to take them to an onboarding login screen, where they enter their corporate username and password. The captive web portal authentication validates the employee's username and password via RADIUS and LDAP. The captive web portal authentication is protected via HTTPS.

After the employee logs in to the network, certificates and/or an onboarding application are downloaded to the mobile device. If an onboarding application is part of the process, the application is executed to securely download the 802.1X/EAP root certificate and/or other security credentials as well as provision the supplicant on the mobile device. Custom onboarding applications are often distributed via the open SSID. An onboarding solution for different types of users might also be available as a web-based application. Figure 3-4 shows a custom onboarding web application used by the Calgary Board of Education used to provision different security credentials to guests, students, and staff.



Source: Calgary Board of Education

FIGURE 3-4: Onboarding application.

After the employee device is provisioned with the necessary certificates, it is ready to connect to the secure network. Because the secure network is a separate SSID, the employee would have to manually disconnect from the open SSID and reconnect to the secure SSID. Furthermore, some onboarding applications may also automatically remove the network profile of the onboarding SSID and select the employee SSID as the preferred network. Once provisioned, the user experience is enhanced because the client device does not inadvertently associate again to the onboarding SSID.

IN THIS CHAPTER

- » Understanding why guest access is important
- » Keeping your guests on the right network
- » Making sure guests are who they say they are with secure logins
- » Setting up captive web portals
- » Providing self-registration access and employee sponsorship

Chapter 4

Welcome, Guest!

Enterprise WLANs keep your regular employees happy by allowing them to use wireless mobility. Your guests need that kind of freedom as well, but you do not want to give them too much freedom — they do not need access to your prized corporate systems. Customers, consultants, vendors, and contractors often just need access to the Internet to accomplish job-related duties. When guests are more productive, employees will also be more productive.

Guest access can also be a value-added service and often breeds customer loyalty. In today's world, business customers have come to expect guest WLAN access. Free guest access is often considered a value-added service. There is a chance that your customers will move toward your competitors if you do not provide guest WLAN access. Retail, restaurants, and hotel chains are all prime examples of environments where wireless Internet access is often expected by customers. The good news is that most NAC solutions now provide guest management capabilities. This chapter takes a look at how you can offer this service with the right level of security to keep everybody online and happy.

Putting Out the Welcome Mat

The primary purpose of a guest WLAN is to provide a wireless gateway to the Internet for company visitors and/or customers. Because guest users do not need access to company network resources, a guest WLAN must protect the company network infrastructure from the guest users. In the early days of Wi-Fi, guest networks were not very common because of fears that the guest users might access corporate resources. Guest access was often provided on a separate infrastructure. Another common strategy was to send all guest traffic to a separate gateway that was different from the Internet gateway for company employees. For example, a company might use a T1 or T3 line for the corporate gateway, whereas all guest traffic was isolated on a separate DSL phone line.

WLAN guest access has grown in popularity over the years, and the various types of WLAN guest solutions evolved to meet the need. At a minimum, guest WLANs should offer a separate guest SSID, a unique guest VLAN, and a guest firewall policy. Captive web portals and guest access options, such as guest self-registration, also offer security and convenience.

Guest SSID

In the past, a common WLAN design strategy was to segment different types of users, even employees, on separate SSIDs mapped to independent VLANs. For example, a hospital might use unique SSID/VLAN pairs for doctors, nurses, technicians, and administrators. That strategy creates a great deal of layer 2 overhead by using multiple SSIDs. Today, the more common method places all employees on the same SSID and leverages *Remote Authentication Dial-In User Service* (RADIUS) attributes to assign different groups of users to different VLANs. (See Chapter 2 for a reminder). However, all guest user traffic should still be segmented onto a separate SSID. The guest SSID always uses different security parameters from the employee SSID. For example, employee SSIDs commonly are protected with 802.1X/EAP security, whereas guest SSIDs are most often an open network with a captive web portal for authentication.



Although encryption is not usually provided for guest users, some WLAN vendors have begun to offer encrypted guest access and provide data privacy using unique per-user or per-device PSK credentials. Encrypted guest access can also be provided with 802.1X/EAP with Hotspot 2.0 for Passpoint-enabled client devices.

Like all SSIDs, a guest SSID should never be hidden and should have a simple name, such as *Aerohive-Guest*. In most cases, the guest SSID is prominently displayed on a sign in the lobby or entrance of the company offices.

Guest VLAN

Guest user traffic should be segmented into a unique VLAN tied to an IP subnet that does not mix with the employee VLANs. Segmenting your guest users into a unique VLAN is a security and management best practice. Security experts can debate about whether the guest VLAN should be supported at the edge of the network. A frequent design isolates the guest VLAN in what is known as a *demilitarized zone* (DMZ). The guest VLAN does not exist at the access layer. Therefore, all guest traffic must be tunneled from the AP back to the DMZ, where the guest VLAN does exist. An IP tunnel, commonly using the *Generic Routing Encapsulation* (GRE) protocol, transports the guest traffic from the edge of the network back to the isolated DMZ. Depending on the WLAN vendor solution, the tunnel destination in the DMZ can be either a WLAN controller, GRE server appliance, or a router.

Isolating the guest VLAN in a DMZ has been a common practice for many years. However, it is no longer necessary if you enforce guest firewall policies at the edge of the network. Various WLAN vendors now build enterprise-class firewalls into access points. If the guest firewall policy can be enforced at the edge of the network, the guest VLAN can also reside at the access layer and no tunneling is needed. Why not make life easier for everybody involved?

Guest Firewall Policy

The most important security component of a guest WLAN is the firewall policy. The guest WLAN firewall policy prevents guest user traffic from getting near the company network infrastructure and

resources. Figure 4-1 shows a very simple guest firewall policy that allows DHCP and DNS but restricts access to private networks 10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16. Guest users are not allowed on these private networks, because corporate network servers and resources often reside on that private IP space. The guest firewall policy should simply route all guest traffic straight to an Internet gateway and away from the corporate network infrastructure.

| Source IP | Destination IP | Service | Action |
|-----------|-------------------------|-------------|--------|
| Any | Any | DHCP-Server | PERMIT |
| Any | Any | DNS | PERMIT |
| Any | 10.0.0.0/255.0.0.0 | Any | DENY |
| Any | 172.16.0.0/255.240.0.0 | Any | DENY |
| Any | 192.168.0.0/255.255.0.0 | Any | DENY |
| Any | Any | Any | PERMIT |

Source: Aerohive

FIGURE 4-1: Guest firewall policy.

Make sure you permit traffic for these firewall ports:

- »» DHCP server (UDP port 67)
- »» DNS (UDP port 53)
- »» HTTP (TCP port 80)
- »» HTTPS (TCP port 443)

These ports allow the guest user's wireless device to receive an IP address, perform DNS queries, and browse the web. Many companies require their employees to use a secure VPN connection when connected to an SSID other than the company SSID. Therefore, it is recommended that you also allow access for IPsec IKE (UDP port 500) and IPsec NAT-T (UDP port 4500).


The guest firewall policy can be much more restrictive. Depending on company policy, many more ports can be blocked. One practice is to force the guest users to use webmail and block SMTP and other email ports so that users cannot "spam through" the guest WLAN. However, now that most mail services use SSL, this

practice is not as common. The security policy of the company should determine which ports need to be blocked on the guest WLAN. If the policy forbids the use of SSH on the guest WLAN, then TCP port 22 will need to be blocked. In addition to blocking UDP and TCP ports, several WLAN vendors now have the capability to block applications. In addition to stateful firewall capability, WLAN vendors have begun to build application-layer firewalls capable of *deep packet inspection* (dpi) into access points or WLAN controllers. An application-layer firewall can block specific applications or groups of applications. For example, some popular video streaming applications can be blocked on the guest SSID. The company security policy also determines which applications should be blocked or rate-limited on a guest WLAN.

Captive Web Portals

Often, guest users must log in through a captive web portal page before they are given access to the Internet. You have probably used a captive web portal when logging on to the Wi-Fi at an airport or hotel. One of the most important aspects of the captive web portal page is the legal disclaimer. A good legal disclaimer informs guest users about acceptable behavior when using the guest WLAN. Businesses are more likely to be legally protected if something bad, such as being infected by a computer virus, should happen to a guest user's WLAN device while connected through the portal. A *captive portal* solution effectively turns a web browser into an authentication service.

To authenticate, the user must first connect to the WLAN and launch a web browser. After the browser is launched and the user attempts to go to a website, no matter what web page the user attempts to browse to, the user is redirected to a different URL, which displays a captive portal login page. Captive portals can redirect unauthenticated users to a login page using an IP redirect, DNS redirection, or redirection by HTTP. As shown in Figure 4-2, many captive web portals are triggered by DNS redirection. The guest user attempts to browse to a web page, but the DNS query redirects the browser to the IP address of the captive web portal. You shall not pass — at least until we know who you are.



Source: Aerohive

FIGURE 4-2: Captive web portal — DNS redirection.

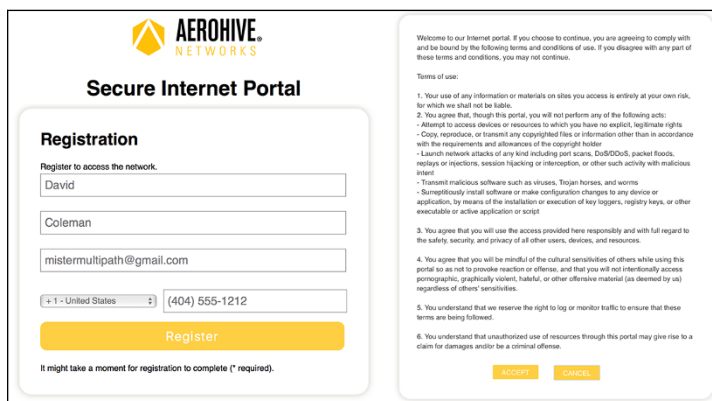
Captive portals are available as a stand-alone server solution or as a cloud-based service. Additionally, most WLAN vendors offer integrated captive portal solutions. The captive portal may exist within a WLAN controller, or it may be deployed at the edge within an access point. Vendors that support captive portals provide the ability to customize the captive portal page. NAC solutions also utilize captive web portal authentication.

You can typically personalize the page by adding graphics, such as a company logo, inserting an acceptable use policy, or configuring the login requirements. Depending on the chosen security of the guest WLAN, different types of captive web portal login pages can be used. A user authentication login page requires the AP or switch to query a RADIUS server with the guest user’s name and password. If the guest user does not already have an account, the login page may provide a link, allowing the user to create a guest account, as shown in Figure 4-3.

The guest registration page allows the user to enter the necessary information to self-register. If identity credentials are not necessary, the guest user may connect to a captive portal web page that only requires acknowledgement of a user policy acceptance agreement. Enterprise businesses can also create different captive web portal experiences for different types of users as well as custom captive web portals for different business locations.

Captive portal authentication often uses RADIUS servers to validate guest user credentials for a guest SSID. A captive web portal solution queries a RADIUS server with a username and password, using a weak authentication protocol such as MS-CHAPv2.

As opposed to using a preexisting user database, such as Active Directory, the guest credentials are usually created during the guest registration process and often stored in the native database of the RADIUS server. Captive web portal authentication is also often used for validating employee credentials for BYOD devices. In that case, the employee database would most likely be Active Directory, which, in turn, would be queried by the RADIUS server.



Source: Aerohive

FIGURE 4-3: Check out these sample captive web portals.

Keep in mind that a captive web portal requires user interaction, and sometimes the guest user experience can be troublesome. Captive web portals often fail after browser updates or mobile device operating system updates. DNS problems also cause captive web portal failures. And, to be honest, the design of many captive web portals is not always user-friendly. At some point in time, just about everyone has had a bad experience with a captive web portal. Make your captive web portals simple, easy to understand, and thoroughly tested to provide the best user experience. You want your guests to be happy, don't you?

Isolate, Limit, and Filter

When guest users are connected to the guest SSID, they are all in the same VLAN and the same IP subnet. Because they reside in the same VLAN, the guests can perform peer-to-peer attacks against

each other. Nobody likes it when guests fight at a dinner party, and nobody likes guests fighting on your network. These features help you keep the peace.

Client isolation

Enable *client isolation* on WLAN access points to block wireless clients from communicating directly with other wireless clients on the same wireless VLAN. This feature prevents packets arriving at the AP's wireless interface from being forwarded back out of the wireless interface to other clients. Isolate each user on the wireless network to ensure that a wireless station cannot be used to gain Layer 3 or higher access to another wireless station. The client isolation feature is usually a configurable setting per SSID linked to a unique VLAN. Client isolation is highly recommended on guest WLANs to prevent peer-to-peer attacks.

Rate limiting

Enterprise WLAN vendors also offer the capability to throttle the bandwidth of user traffic. *Rate limiting* (also known as bandwidth throttling) curbs traffic at either the SSID level or the user level. This feature ensures that the majority of the bandwidth is preserved for employees. Rate limiting the guest user traffic to 1,024 Kbps is a common practice. However, because guest access is usually an expected value-added service, rate-limiting on guest SSIDs may not be a good strategy. Some businesses that attempt to monetize WLAN guest access often provide two levels of guest access. The free level of guest access is rate-limited, whereas the paid guest access has no bandwidth restrictions.

Web-content filters

Enterprise companies often deploy *web-content filters* to restrict the type of websites that their employees can view while at the workplace. A web-content filtering solution blocks employees from viewing websites based on content categories. Each category contains websites or web pages that have been assigned based on their primary web content. For example, the company might use a web-content filter to block employees from viewing any websites that pertain to gambling or violence. Content filtering is most often used to block what employees can view on the Internet,

but web-content filtering can also be used to block certain types of websites from guest users. All guest traffic might be routed through the company's web-content filter.

Guest Management

As Wi-Fi evolved, so has WLAN guest-management solutions. Most guest WLANs require guest users to authenticate with credentials via a captive web portal. Therefore, you must create a database of user credentials. Unlike user accounts in a preexisting Active Directory database, guest user accounts are normally created on the fly in a separate guest user database.

Guest user information is usually collected when the guests arrive at company offices. Someone has to be in charge of managing the database and creating the guest user accounts. IT administrators are typically too busy to manage a guest database; therefore, the individual who manages the database is often a receptionist or the person who greets guests at the front door. This individual requires an administrative account to the guest-management solution, which might be a RADIUS server or some type of other guest database server. The guest-management administrators have the access rights to create guest user accounts in the guest database and issue the guest credentials, which are usually usernames and passwords.

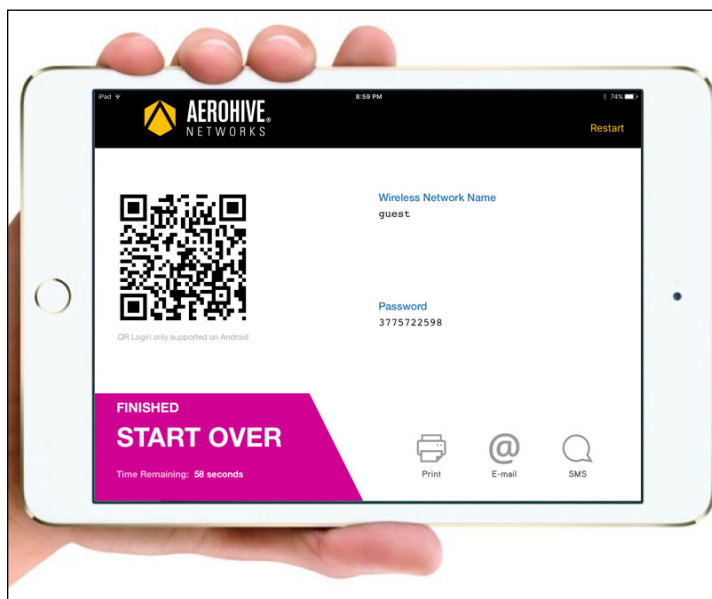


REMEMBER

A guest-management solution can be cloud-based or reside as an on-premise NAC server in the company datacenter. Although most guest-management systems are built around a RADIUS server, the guest-management solution offers features in addition to providing RADIUS services. Modern WLAN guest-management solutions offer robust report-generation capabilities for auditing and compliance requirements. A guest-management solution can also be used as a 24/7 full-time monitoring solution. An IT administrator usually configures the guest-management solution initially. However, a company receptionist will have limited access rights to provision guest users. NAC solutions with guest-management capabilities can also be integrated with LDAP for employee sponsorships and usually have some method for guest users to self-register. Most often, guest-management solutions

are used for wireless guests, but they might also be used to authenticate guests connected to wired ports.

As you can see in Figure 4-4, you can deliver the guest credentials to the guest user in multiple ways. The credentials can be delivered via an electronic wallet, QR code, SMS text message, an email message, or a printed receipt. The SMS, email, and receipt can also be customized with company information. The guest registration login pages can all be customized with company logos and information.



Source: Aerohive

FIGURE 4-4: Guest credential delivery methods.

Mystery Guest, Self-Register Please!

NAC and guest-management solutions traditionally relied on a company receptionist or lobby ambassador to register the guest users. A good guest-management solution allows the receptionist to register a single guest user or groups of users. Over the past few years, more companies allow guest users to create their

own account, via a *self-registration* process. When guest users are redirected to the captive web portal, a link on the login web page redirects unregistered users to a self-registration page. Simple self-registration pages allow guests to fill out a form, and the system creates a guest account. The guest either receives a display or printout of the necessary credentials. More advanced self-registration pages require guests to enter an email or SMS address, which the registration system uses to send users their login credentials.

As shown in Figure 4-5, some guest-management solutions now offer kiosk applications, where the self-registration login page runs on a tablet that functions as the kiosk.



Source: CloudCarrier

FIGURE 4-5: Guest access kiosk.



TIP

Deploy the kiosk in the main lobby or at the entrance to the company. That way, the receptionist does not have to assist the users and can concentrate on other duties.

Can You Vouch for Me?

Guest users might also be required to enter the email address of an employee, who in turn must approve and sponsor the guest. The sponsor typically receives an email with a link that allows them to easily accept or reject the guest's request. Once the user is registered or sponsored, they can log in using their newly created credentials. A guest-management solution with *employee sponsorship* capabilities can be integrated with an LDAP database, such as Active Directory.

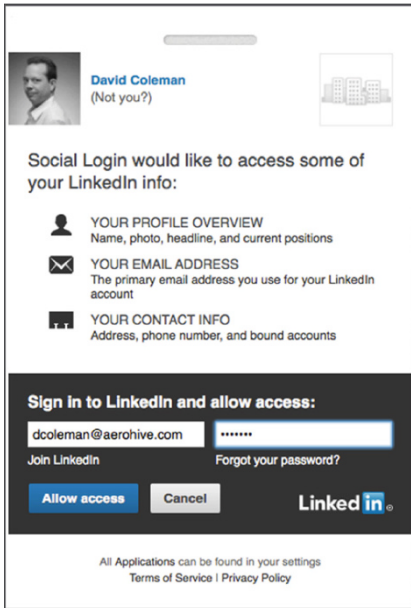
A receptionist can register guest users, or a company may choose to use a registration kiosk so that guests can self-register. Self-registration with employee sponsorship is becoming popular for many organizations.

When guest users initially connect to the guest network, they are redirected to a captive portal page. The captive portal page prompts them to log in if they already have an account, or it allows them to click a link that allows them to create their own guest account. The guest must enter the email address of the employee who is sponsoring them. Typically, the guest has a business meeting with the employee who is providing sponsorship (or at least a quick cup of coffee).

When the registration form is completed and submitted, the sponsor receives an email notifying them that the guest would like network access. The email typically contains a link that the sponsor must click to approve network access. Once the sponsor clicks the link, the guest account is approved. The guest then receives confirmation by email or SMS, and they will be allowed to log in to the network. If the sponsor does not click the link, the guest account is never created, and the guest is denied access to the network. Employee sponsorship ensures that only authorized guest users are allowed onto the guest WLAN and that the company employees are actively involved in the guest user authorization process.

Social Login

Social networks can work in business, too! *Social login* is a new trend in guest networks in retail and service industries. This method uses existing login credentials from a social networking service, such as Twitter, Facebook, or LinkedIn, to register on a third-party website. Social login allows a user to forgo the process of creating new registration credentials for the third-party website. Social login is often enabled using the *Open Standard for Authorization (OAuth)* protocol. OAuth 2.0 is a secure authorization protocol that allows access tokens to be issued to third-party clients by an authorization server. As shown in Figure 4-6, the OAuth 2.0 authorization framework enables a third-party application to obtain limited access to an HTTP service and can be used for social login for Wi-Fi guest networks.



Source: Aerohive

FIGURE 4-6: OAuth 2.0 application.

Social login can be tied to an open guest SSID. Guest users are redirected to a captive web portal page, where they can then log in to the guest WLAN using their existing social media login credentials. Retail and service businesses like the idea of social login because it allows them to obtain meaningful marketing information about guest users from the social networking services. Businesses can then build a database of the type of customers who are using the guest Wi-Fi while shopping at the business.



WARNING

There are serious privacy concerns with social login, and the login captive web portal always has a legal disclaimer stating that customer information might be gathered if the customer agrees to use the social login registration to the guest WLAN.

SAML and SSO

Many NAC solutions now offer support for the *Security Assertion Markup Language* (SAML) for single sign-on (SSO) capabilities. SAML is an open-source language for exchanging user authentication data as XML between trusted parties. For example, a Google account user may utilize SAML to login to multiple cloud applications via SSO. *Identity and access management* (IAM) service providers such as OKTA or Ping Identity provide administrators with a central location to manage all users and cloud applications. The IAM service provider enables users with *single sign-on* (SSO) capabilities for all cloud applications using SAML. Another growing networking trend is to provide guest access via a captive web portal using SAML. Guest users use their existing SSO credentials to authenticate via the guest portal. Most enterprise-grade NAC solutions now offer support for SAML and SSO for guest access and other authentication scenarios. IAM service providers may also use SAML to enable *multifactor authentication* (MFA).

Encrypted Guest Access

Most guest networks are open networks that do not use encryption. Thus, there is no data privacy for guest users. Numerous wireless attacks exist that make unsecured Wi-Fi users vulnerable. Because most guest WLANs do not use encryption, guest

users are low-hanging fruit and often targets of skilled hackers or attackers. This is why so many people warn you not to use networks in airports or hotels.

For that reason, many corporations require their employees to use an IPsec VPN solution when connected to any kind of public or open guest SSID. Because the guest SSID does not provide data protection, guest users must bring their own security in the form of a VPN connection that provides encryption and data privacy.

The problem is that many consumers and guest users are not savvy enough to know how to use a VPN solution when connected to an open guest WLAN. As a result, recent trends point to providing encryption and better authentication security for WLAN guest users. Protecting the company network infrastructure from attacks from a guest user remains a top security priority. However, if a company can also provide encryption on the guest SSID, the protection provided to guest users is a value-added service.

One simple way to provide encryption on a guest SSID is to use a static PSK. Although encryption is provided when using a static PSK, this is not ideal because the static PSK is susceptible to social engineering attacks. Some WLAN vendors offer cloud-based solutions to distribute secure guest credentials in the form of unique per-user PSKs. Aerohive's *Private Pre-Shared Key* (PPSK) solution is often used for guest access security to provide unique identity credentials as well as data privacy with encryption for guest users.



TECHNICAL
STUFF

A more complex method, sometimes used by telecom service providers, offers encrypted guest access via 802.1X/EAP for Passpoint-enabled client devices. Passpoint is the brand for the certification program operated by the Wi-Fi Alliance based on the HotSpot 2.0 specification.

- » Understanding posture assessment
- » Utilizing RADIUS CoA
- » Automating device access dynamically

Chapter 5

You've Got a NAC for This!

Part of keeping your network secure obviously involves restricting which devices access your network. But you cannot just give full access to every device on your network, either. Just as you do not allow party guests into every room in your house (the cat just gets nervous around new people), you do not let every authenticated device run wild on your network. In this chapter, you learn how to manage access to your network and give devices exactly the level of access they need.

Take Control

Network access control (NAC) evaluates the capability or state of a computer and determines the level of access to allow. NAC has evolved over the years from a system that primarily assessed the virus and spyware health risk to an environment where checks and fingerprinting are performed on an endpoint, extensively identifying its capabilities and configuration. These checks are integrated with 802.1X/EAP and RADIUS to authenticate and authorize network access for the user and the computer.

NAC solutions safely enable appropriate network access for diverse endpoints. These actions may include preventing laptops that lack antivirus, patches, or host intrusion prevention software from accessing the network and placing other computers at risk of cross-contamination from viruses and other malware. Think of this as a system for telling your coworkers to stay home when they have a cold (and what a wonderful world that would be). Where conventional IP networks enforce access policies in terms of IP addresses, NAC environments allow a network administrator to grant access based on user or group membership, in combination with the current posture of the endpoint.



TIP

NAC solutions allow network operators to define policies, such as the types of computers or roles of users allowed to access areas of the network, and enforce them in access points, switches, routers, firewalls, and other access network infrastructure.

Check Your Posture

NAC began as a response to computer viruses, worms, and malware that appeared in the early 2000s. The early NAC products date back to around 2003 and provided what is known as posture assessment. *Posture assessment* applies a set of rules to check the health and configuration of a computer and determine whether it should be allowed access to the network. NAC solutions do not perform the health checks themselves. Instead, NAC solutions validate that a device adheres to policy. Posture assessment verifies that security software, such as antivirus, antispymware, and a firewall, is installed, up-to-date, and operational. Figure 5-1 shows an example of reported violations based on a posture assessment scan.



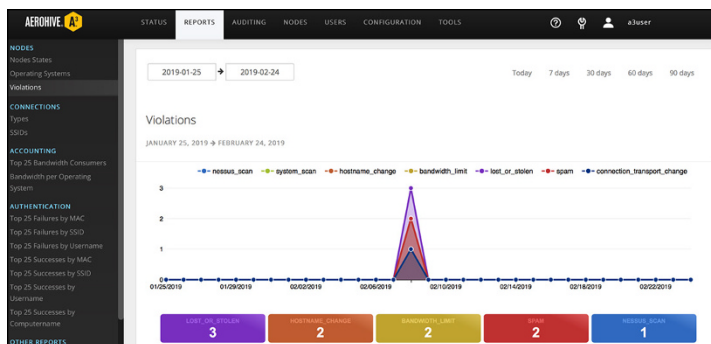
REMEMBER

Essentially, posture assessment “checks the checkers.” In addition to checking security software status, posture assessment can check the state of the operating system. Security policy can make sure that specific patches or updates are installed, verify that certain processes are running or not running, or even check to determine whether or not specific hardware (such as USB ports) is active.

Talk to my agent

Posture checking uses many methods. One popular approach utilizes a *persistent agent*, which is software that is permanently installed on the computer. The agent is often capable of making

real-time assessments that detect changes in posture immediately. When a company deploys a posture-checking software, the persistent agent is likely installed on all the corporate laptops to make sure that they are healthy. Think of it like introducing herd immunity to the digital world.



Source: Aerohive

FIGURE 5-1: A posture exam. . . I mean, scan.

A *dissolvable agent* is software that is temporarily installed on an endpoint during the initial network access. This tool may be useful for scanning guest computers, because guests are unlikely to allow your company to permanently install software on their laptop. Some solutions use network scanning technologies to assess the posture of the endpoint. For example, when a user connects to the network, a Windows Management Instrumentation (WMI), OpenVAS, or Nessus scan may launch to assess the status of that computer before access decisions are determined. Subsequent scans may run on a configurable interval. Lastly, many NAC solutions act as the enforcement mechanisms for other solutions' posture assessments. For example, if an MDM solution deems an endpoint is out of compliance with its policy, it can trigger an event in the NAC to restrict access.

The end-user experience is a very important aspect of successful posture assessment. If a computer is found to be out of compliance with policy, several actions are possible. An administrator may elect to merely alert on the issue in the NAC dashboard, with no user interaction. Another option is to notify the end user of the issue via a message in their web browser or email, sharing instructions for solving the issue and potentially issuing warnings

and resolution requirements. Further, the administrator may opt to quarantine the computer until the issue is resolved.

Generally, quarantined endpoints have access to the network resources required to resolve the problem (such as installing an antivirus client) but no additional access. The ideal scenario would allow the persistent agent to automatically fix or remediate the problem so that the computer can pass the check and gain network access. Because the persistent agent is installed on the corporate computer and typically has permissions to make changes, that software can perform automatic remediation.



TIP

Automatic remediation is often complex to configure, and many posture assessment solutions do not include this capability, relying instead on messaging to the end user for how to resolve the issue themselves. Self-remediation is often a better choice.

Back on the path

Providing a path back to network access is important for users who have been quarantined for noncompliance with security policy. You cannot just strand users and not help them fix the problem. Often, this involves a communication plan about the reasons for implementing the requirements, and a process for regaining network access. Some NAC solutions immediately restore network access on resolution of the issue; other solutions may require the user to reconnect to the network for a new assessment.

RADIUS Change of Authorization

RADIUS Change of Authorization (CoA) is a feature that allows a RADIUS server to adjust an active client session. CoA messages are used by the AAA framework to dynamically modify subscriber sessions. For example, if the posture of the endpoint of an endpoint changes, a RADIUS CoA message can update the authorization policy of the endpoint.

Prior to RADIUS Change of Authorization (CoA), if a client were authenticated and assigned a set of permissions on the network, the client authorization would not change until the client logged out and logged back in. This setup only allowed an authorization decision during the initial connection of the client.

To use an analogy to explain RADIUS CoA, let's say Jack is going to a club with friends to enjoy some cocktails and dancing. When they arrive, a bouncer at the door admits them into the club but tells them that they are not allowed to become drunk or cause trouble in the club. While telling them this, the bouncer checks to make sure that they are not already drunk or causing trouble. Unfortunately, the bouncer must stand at the door and monitor the guests only as they enter the club. The bouncer cannot monitor the guests once they are inside the club. After a few nights of experiencing some problems in the club, the manager decides to hire additional security personnel, who walk around the club and monitor guests already in the club. Anyone who is found to be drunk or causing problems is either restricted within the club (maybe they are no longer allowed to purchase alcoholic drinks), or removed from the club. Once such a guest is outside the club, the bouncer at the door can reevaluate the status of the guest, possibly denying reentry into the club, allowing the guest back in the club, or allowing the guest to reenter but with a different set of permissions.



RADIUS CoA was originally defined by RFC 3576 and later updated in RFC 5176. Many of the AAA servers, NAC servers, and enterprise wireless equipment reference RADIUS RFC 3576 on configuration menus without referring to CoA. Therefore, from a practical perspective, you should be aware that if you see RFC 3576 on any configuration menu, that is the section where RADIUS CoA is configured.

Be Dynamic

NACs also automatically and dynamically determine network access for a given device. This access could be based on an initial assessment upon connecting to the network, or a change in access afterwards based on a change in the endpoint's network behavior or status.

RADIUS accounting (the final A in AAA) monitors the user connection. In the early days of AAA, RADIUS accounting typically tracked client connection activity. In some environments, logging in and logging off events may be all you want or need to track. Enhancements to accounting allow the AAA server to also provide interim accounting. Interim accounting can track resource

activity, such as time and bytes used for the connection. As shown in Figure 5-2, if a user exceeds or violates the allowed limits of resources, RADIUS CoA can dynamically change the permissions that the user has on the network. Changes in endpoint behavior can also trigger RADIUS CoA and therefore a change in network access rights. For example, a NAC may transition an endpoint to the isolation role and isolation VLAN if malware is detected. Additionally, a bandwidth violation may result in the endpoint device and user moving back into the registration role.

NAC solutions can also send identity and session information to other infrastructure devices for their use in decision making. For example, username to IP information can be published to a firewall so that security policies at the edge of the network can be based on group membership rather than source IP. This configuration is especially helpful in diverse networks, where not all endpoints are managed by group policy.

| Id | Description | Actions | Target Role |
|---------|-----------------|-----------------------------------------------------------------------------------------|--------------|
| 2000000 | Malware | Change network access on violation Email violation to administrator Log violation | isolation |
| 2000032 | LSASS Exploit | Email violation to administrator Log violation | isolation |
| 2001569 | NetBIOS Scan | Email violation to administrator Log violation | isolation |
| 2001904 | Telnet Scan | Email violation to administrator Log violation | isolation |
| 1200002 | Time Expiration | Change network access on violation Log violation | registration |
| 1200003 | Bandwidth Limit | Change network access on violation Log violation | registration |

Source: Aerohive

FIGURE 5-2: CoA: You have been demoted!

IN THIS CHAPTER

- » Identifying the type and OS for devices connecting to your network
- » Getting past spoofed information
- » Determining network access via device fingerprinting

Chapter 6

You Have a Great Profile!

One of your best tools in managing the security of your network is learning exactly what devices are trying to connect to your network and what software they're running. If an out-of-date device with an older operating system tries to infiltrate your wireless network and possibly leave gaping holes, you will want to stop it before anything gets started. In this chapter, you learn about methods used to conduct this process, often referred to as either *device profiling* or *device fingerprinting*. The operating system of both wireless and wired client devices can be determined by a variety of fingerprinting methods, including utilizing a posture assessment agent, OUI matching, DHCP analysis, and HTTP snooping. Sounds like a detective novel, doesn't it? Device fingerprinting provides network administrators visibility as to what type of devices are connecting to the corporate network as well as controlling secure access to certain types of devices.

Talk to My Agent

The most accurate perspective on the type of device and OS comes when a posture assessment agent is installed on an endpoint. Posture assessment agents offer increased access to a device, and therefore may be able to distinguish between OS versions and patch levels. While this is useful for corporate-owned devices

such as laptops, agents are not commonly installed on the BYOD and IoT endpoints, where there is greater diversity and a real need for visibility. For more on posture assessment, see Chapter 5.

The *organizationally unique identifier* (OUI) address is the first three octets of a device MAC address that identifies the manufacturer of the networking device. The remaining octets of the MAC address are unique and are used to identify the individual device. The IEEE assigns OUI addresses to manufacturers. Most solutions utilize OUI matching as a first-level of device fingerprinting. For example, OUI matching is a simple method of distinguishing devices manufactured by Apple from those produced by Samsung.



REMEMBER

DHCP analysis utilizes the DHCP handshake to determine the type of device that is requesting a network address. After a client successfully establishes a layer 2 connection, the next action is to send a DHCP request to obtain an IP address. As part of that request, the client device includes DHCP option information and requests a list of DHCP parameters or options from the DHCP server. These options may include subnet mask, domain name, default gateway, and the like.

When a client sends DHCP discover and request messages, each type of client requests different parameters under the DHCP option 55 portion of the request. The parameters within DHCP option 55 create a fingerprint that can be used to identify the operating system of the client.

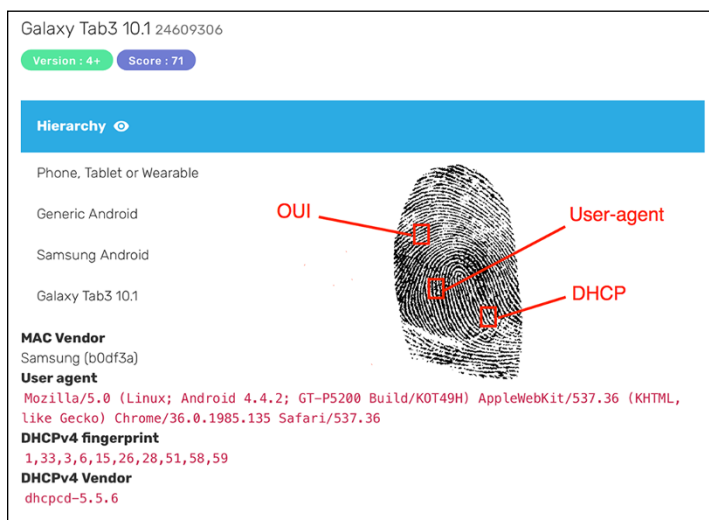
For example, iOS devices include a common set of parameters when performing a DHCP request, thus making it possible to identify that the device is most likely an iOS device. DHCP fingerprinting is not perfect, and it is often not possible to discern the difference between similar devices, such as an iPod, an iPhone, or an iPad.



TECHNICAL
STUFF

Depending on the NAC vendor, the DHCP fingerprint is referenced as an ASCII list of parameter request options, such as 1, 3, 6, 15, 119, 252. Or it might be displayed as a hexadecimal string, such as 370103060F77FC. In the string, the first two hex digits are equal to ASCII 55 (option 55), and each of the following two digits pairs are the hex values of each option. You can find an extensive list of DHCP fingerprints at www.fingerbank.org. Although the parameter request list is not guaranteed to be unique, it can typically be used along with other fingerprinting techniques to identify devices.

Another OS detection method is *HTTP fingerprinting*. The user-agent header within an HTTP packet identifies the client operating system. During captive portal authentication, NAC solutions are able to inspect HTTP and HTTPS frames while handling the client requests. This fingerprinted information is combined with the information obtained through other methods to paint a better picture of the client device. Figure 6-1 shows the device profile of a Samsung Galaxy tablet using the Android OS identified by the OUI address, DHCP fingerprint, and the HTTP user-agent fingerprint.



Source: Aerohive

FIGURE 6-1: Fingerprints make good impressions.

Other ways of obtaining client OS information are methods such as SNMP and TCP scanning. For example, the initial *time to live* (TTL) in the IP header and the TCP receive window size of the first packet in a TCP session can also be used to identify some operating systems.



An end user who is determined to circumvent network access restrictions can misrepresent their device type. For example, it is possible to edit the user-agent string sent by the browser or spoof the MAC address associated with the network adapter. Sophisticated NAC solutions use a combination of distinct methods to fingerprint an endpoint, with confidence ratings in the results

and processes for resolving discrepancies between fingerprinting methods. Additionally, a robust NAC offers the capability to create custom device classification for endpoint devices. This step is particularly important to ensure that unknown IoT devices are properly classified for visibility and security enforcement.



REMEMBER

Keep in mind that the best way to ensure user compliance with access requirements is a path of least resistance for the user. If users have an easy way to onboard their gaming device, they are less likely to edit the device user-agent strings to appear as a Windows OS device.

Using Device Fingerprinting to Determine Network Access

Once an administrator has accurate information about the types of devices in use on the network, they can establish differentiated access based on that data. For example, an organization may want to allow laptops, tablets, and smartphones on its guest network and may elect to prohibit gaming systems and other media devices. Another common use case defines a list of allowed operating systems on their corporate or secure SSID, to protect that portion of the network from nonapproved OS versions and devices.



TIP

IoT devices, such as sensors, may be placed in a very specific VLAN with very specific access policies. For example, patient monitoring equipment in a hospital may connect to the corporate SSID but be restricted to a specific VLAN and only have access to the necessary monitoring servers on the hospital network. The explosion of IoT devices in the workplace has made device profiling a clear necessity.

- » Reviewing the functionality of Aerohive A³
- » Managing security in the cloud
- » Reducing the complexity of your security implementation

Chapter 7

Introducing Aerohive A³ — Cloud-Managed NAC

In this chapter, I introduce you to A³, Aerohive's cloud-managed NAC solution. A³ provides an enterprise-grade set of features for device onboarding, visibility, profiling, and policy enforcement. A³ effectively secures all devices on the network — wireless and wired — including corporate devices, BYOD, guest, and IoT.

A³ Functionality

A³ provides an extensive feature set for secure network access and control.

Device onboarding for guest and corporate devices

With A³, administrators have access to a wide variety of tools, including highly customizable captive web portals (CWPs), social login, Hotspot 2.0, and eduroam for higher education institutions.

Authentication methods for guest and corporate onboarding

Supported options include 802.1X certificates, pre-shared keys (PSKs), and multi-factor authentication. A³ integrates with LDAP and Active Directory. A³ also supports OAuth2 and integrates with SAML for single sign-on (SSO) functionality.

Once onboarded, devices can be provisioned using the following methods:

- » Agents for Android and Windows devices
- » Apple macOS and iOS devices via the device API
- » Integration with all leading mobile device management (MDM) and enterprise mobility management (EMM) solutions.

Network access control (NAC)

A³ supports both policy-based and compliance-based network access control.

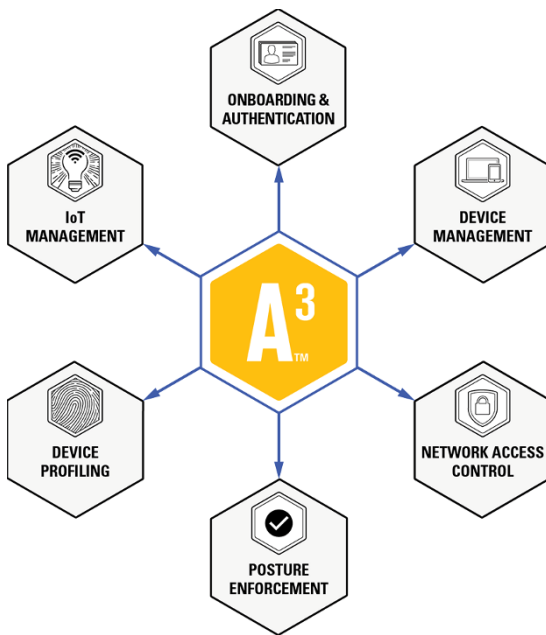
- » Policy-based access control is the first line of defense against unauthorized network access. A³ enables granular, identity-based policies that determine what resources users and devices can access, depending on their role.
- » Compliance-based network access control is used to ensure that devices stay secure over time and cannot spread acquired vulnerabilities (for example, viruses) into the network.

Figure 7-1 provides an overview of A³ functionality.

As part of the network access control (NAC) functionality, A³ also offers extensive reporting and accounting information.

Device profiling

A³ includes a large database of electronic device signatures, or fingerprints. This repository allows A³ to automatically recognize the make and type of a device, differentiating, for example, between an Apple MacBook Air, a Samsung Galaxy S9, or a Nest smart thermostat.



Source: Aerohive

FIGURE 7-1: A starry view of all the available functionality.

IoT security

A³ can also provide highly effective security for IoT devices. Using device profiling capabilities, A³ can automatically identify IoT devices and then onboard them by assigning dedicated VLANs with appropriate, narrowly defined network policies tailored for IoT.

Security orchestration

A³ integrates with all leading firewalls, MDM/EMM solutions, posture assessment solutions, and intrusion detection systems (IDSs). A³ can then exchange information with them and use that information to dynamically adjust network access rights and enforce policies through the integrated solutions, which further strengthens overall network security.

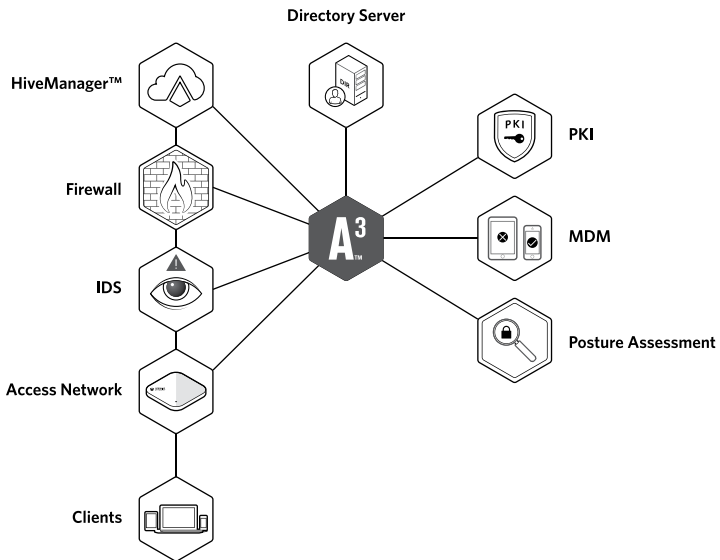
Managed service capabilities

A³ supports additional capabilities for managed service providers (MSPs). This includes federated multitenancy that allows security

service providers to seamlessly manage A³ as a service for multiple customers. A³ also supports paid Internet services for public venues and hotels, thanks to integration with several popular billing providers.

Cloud Management

A³ uses Aerohive's leading cloud networking infrastructure to provide an industry-first capability: *cloud management* of local A³ enforcement nodes. A³ supports cloud monitoring of local A³ instances, as well as cloud-configuration workflows for different types of devices. The cloud centralizes monitoring and configuration of remote sites, while localized tasks like device authentication and access control enforcement are executed by onsite A³ enforcement nodes. This allows A³ customers to enjoy unprecedented levels of deployment flexibility and efficiency. As shown in Figure 7-2, customers who run A³ on an Aerohive access network can leverage cloud management to access A³ via the HiveManager network management system, for single-pane-of-glass visibility.



Source: Aerohive

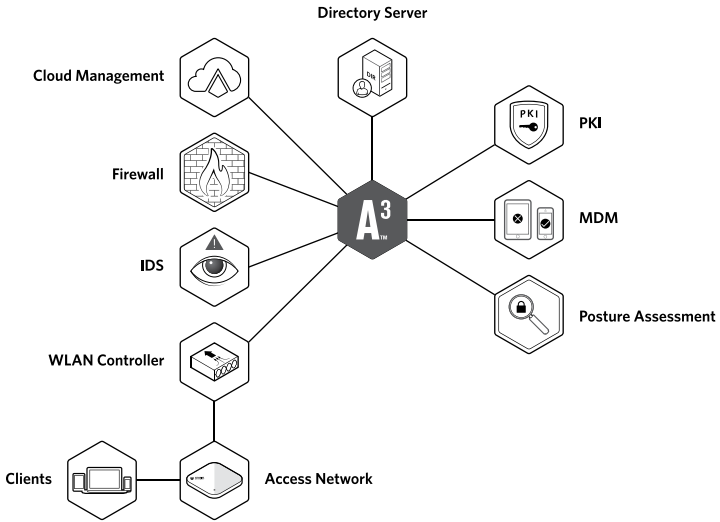
FIGURE 7-2: Keep your head (and security management) in the clouds.



TIP

Customers who prefer an on-premises deployment model can continue to manage A³ on-premises.

The cloud management option of A³ onsite enforcement nodes is also available when A³ is deployed with networks from other vendors. Figure 7-3 shows A³ is network agnostic and works with switches and WLAN equipment from all major vendors.



Source: Aerohive

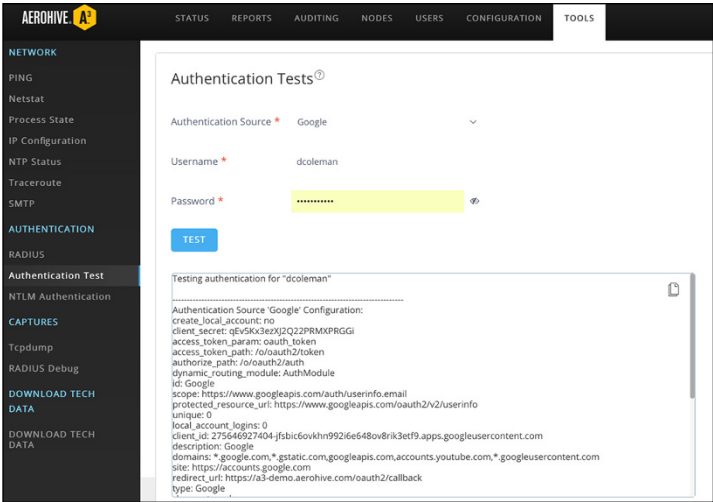
FIGURE 7-3: A³ plays nicely with others.

Streamlining Operations

The centralized monitoring and configuration of the cloud management option significantly streamlines operations. It increases visibility and eliminates the need for repetitive onsite configuration and troubleshooting.

In addition, A³ simplifies complex deployment and management tasks that previously were only accessible through the command line, by incorporating them into the user interface. For example, A³ supports completely UI-based deployments of A³ instances and clusters, allowing the setup of an entire cluster in as little as 30 minutes.

As shown in Figure 7-4, Aerohive A³ also incorporates troubleshooting tools into the UI. These capabilities significantly reduce time, effort, and the need for training when compared with traditional workflows.



Source: Aerohive

FIGURE 7-4: A³ makes it easier to deal with your troubles.

Chapter 8

Ten Things You Need to Know about Cloud-Managed NAC

Want a quick refresher on important points from this book? Need to reinforce some learning from other chapters? Did you just skip to the back of the book to see how it all ended? In all these cases, I have you covered. Keep these ten things in mind when researching your cloud-managed NAC solution, and remember that the butler did it.

Understand the Security Concept of AAA

Authentication, authorization, and accounting (AAA) is a key computer security concept that defines the protection of network resources. Authentication is the verification of identity and credentials. Authorization determines if the device or user is authorized to have access to network resources. Accounting is tracking the use of network resources by users and devices. These concepts help you secure your network, but they will not jumpstart your car.

NAC Is Built on an 802.1X Framework

The 802.1X authorization framework consists of three main components, each with a specific role.

- » **Supplicant:** The endpoint device that asks for access
- » **Authenticator:** The gateway used during authentication, usually a switch or an access point
- » **Authentication server:** The server that validates the supplicant credentials, usually a RADIUS server

These three 802.1X components work together to make sure only properly validated users and devices are authorized to access network resources. A Layer 2 authentication protocol called *Extensible Authentication Protocol (EAP)* is used within the 802.1X framework to validate users at Layer 2.

Modern-Day NACs Provide Certificate Onboarding

One of the biggest challenges for an 802.1X/EAP solution is distributing and installing certificates onto client devices. A device onboarding solution provides a self-service method for an employee to configure a BYOD supplicant and install security credentials, such as an 802.1X root CA certificate and client certificates.

Access Control Is Vital

Role-based access control (RBAC) is an approach to restricting system access to authorized users. The system uses RADIUS attributes to make sure users get the right VLAN, bandwidth allocation, and much more. Additionally, RADIUS *Change of Authorization (CoA)* can transition access dynamically based on a change in the endpoint's network behavior or status. Think of it like telling a drunk laptop to go home. To assist in policy enforcement, most NAC solutions fully integrate with switches, firewalls, *mobile device management (MDM)* solutions, *intrusion detection systems (IDS)*, and third-party endpoint security applications.

Posture Assessment Ensures Endpoint Integrity

This feature does not make sure you will stand up straight, but it does help keep the updates on the up-and-up. Posture assessment ensures endpoint integrity by validating up-to-date versions of device OS, antivirus, antimalware signatures, and application patches. Endpoint devices that do not comply get quarantined until they feel better — that is, they receive the proper updates. Posture assessment checks of endpoints can occur both before and after a device gains network access.

Guest Management Is a Key Component

Even if you are not a coffee shop, everybody expects free Wi-Fi. NAC security requires a means to provide wired or wireless guest access while still protecting the corporate network. Guest management options include social login, captive web portals, self-registration, and employee sponsorship. At a minimum, guest access should require a separate guest SSID, a unique guest VLAN, and a guest firewall policy.

We Live in a BYOD World

Unless you are willing to give everybody free technology, people will want to use their own devices. The *bring your own device* (BYOD) concept allows employees to access the corporate network with personal devices, such as smartphones and tablets. A BYOD policy dictates which corporate resources can or cannot be accessed when employees connect to the company network with their personal devices. A modern-day NAC solution must provide access based on predefined policies for BYOD endpoints in addition to company-owned devices.

Device Identity Is Crucial

Device fingerprinting provides network administrators visibility as to what type of devices are connecting to the corporate network as well as controlling secure access to certain types of devices. Endpoint devices can be identified and classified by a variety of methods, including OUI matching, DHCP analysis, and HTTP snooping. Nightclubs demand identification and sniff out the fake IDs, so think of this feature as your virtual doorman. No snarky velvet rope required!

IoT Is the Next Challenging Frontier for NAC

When everything can access a network, you need a better network to handle it all. A population explosion of IoT devices requiring access to enterprise networks has already begun. A good NAC solution with device fingerprinting capabilities is necessary to identify and classify IoT devices so that proper access policies can be enforced.

Cloud-Managed NAC Is the Future

Next-generation NAC solutions, such as Aerohive's A³, move NAC to the next level with cloud management. A cloud-managed NAC solution scales globally by providing a central point of management for monitoring NAC enforcement nodes at multiple locations. A cloud-managed NAC centralizes monitoring and configuration of remote sites, while localized tasks like device authentication and access control enforcement are executed onsite.

Acronyms

AAA: authentication, authorization, and accounting

AD: Active Directory

API: application programming interface

AVP: attribute-value pair

BYOD: bring your own device

CA: certificate authority

CoA: Change of Authorization

CWP: captive web portal

DHCP: Dynamic Host Configuration Protocol

DMZ: demilitarized zone

DNS: Domain Name System

DPI: deep packet inspection

EAP: Extensible Authentication Protocol

EMM: enterprise mobility management

GPO: Group Policy Object

GRE: Generic Routing Encapsulation

HTTP: HyperText Transfer Protocol

HTTPS: HyperText Transfer Protocol Secure

IAM: identity access management

IANA: Internet Assigned Number Authority

IDS: intrusion detection system

IoT: Internet of Things

LDAP: Lightweight Directory Access Protocol

MAC: medium access control

MDM: mobile device management

MSP: managed service provider

NAC: network access control

OAuth: Open Standard for Authorization

OS: operating system

OUI: organizationally unique identifier

PEAP: Protected Extensible Authentication Protocol

PCI: payment card industry

PKI: public key infrastructure

PPSK: Private Pre-Shared Key

PSK: Pre-Shared Key

RADIUS: Remote Authentication Dial-in User Service

RadSec: RADIUS over TLS

RBAC: role-based access control

SAML: Security Assertion Markup Language

SNMP: Simple Network Management Protocol

SSL: Secure Sockets Layer

SSO: single sign-on

TCP: Transmission Control Protocol

TLS: Transport Layer Security

TTL: time to live

UDP: User Datagram Protocol

VLAN: virtual LAN

VSA: vendor-specific attribute

WLAN: wireless local area network

SECURE YOUR ACCESS NETWORK



Introducing Aerohive A³ – The cloud-managed network access control (NAC) solution that secures all devices on the network, including wireless and wired clients, IoT, and BYOD.

A³ BENEFITS

- Enable streamlined onboarding for thousands of devices
- Identify and secure thousands of IoT devices
- Deploy policy-based and compliance-based access control
- Enjoy the flexibility and scalability of cloud management
- Deploy on any major vendor's network
- Orchestrate complementary security solutions with A³

Learn more at aerohive.com/products/a3/.



NAC with the power of cloud-management

The access network presents a multitude of IT security challenges. An ever-growing number of corporate and guest devices, as well as IoT and BYOD, need to be securely onboarded and provisioned with the access rights appropriate for their roles. A cloud-managed network access control (NAC) solution delivers enterprise-grade functionality for device onboarding, visibility, profiling, and policy enforcement of wired and wireless devices, with the deployment simplicity and scalability of the cloud.

Inside...

- Secure access for employees and guests
- Enable streamlined device onboarding
- Understand how AAA secures your network
- Choose secure authentication methods
- Identify and secure IoT and BYOD devices
- Define role-based access control
- Aerohive A³ — Cloud-Managed NAC



David Coleman is the Senior Technical Evangelist at Aerohive Networks and the co-author of numerous books about Wi-Fi networking and IT security.

Go to **Dummies.com**[®]
for videos, step-by-step photos,
how-to articles, or to shop!

ISBN: 978-1-119-56827-8
Not For Resale

for
dummies[®]
A Wiley Brand



Also available
as an e-book



WILEY END USER LICENSE AGREEMENT

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