

LEARNING MADE EASY

Extreme Networks Special Edition

Wi-Fi 6 & 6E

for
dummies[®]
A Wiley Brand



Enhanced 2.4 and
5 GHz spectral efficiency

Uplink and downlink
multi-user connectivity

Wi-Fi's new home in the
6 GHz frequency band

Compliments
of



ADVANCE WITH US

David Coleman – CWNE #4

About Extreme Networks

Extreme Networks, Inc. (EXTR) creates effortless networking experiences that enable all of us to advance. We push the boundaries of technology leveraging the powers of machine learning, artificial intelligence, analytics, and automation. Over 50,000 customers globally trust our end-to-end, cloud-driven networking solutions and rely on our top-rated services and support to accelerate their digital transformation efforts and deliver progress like never before. For more information, visit www.extremenetworks.com or follow us on LinkedIn, YouTube, Twitter, Facebook, and Instagram.



Wi-Fi 6 & 6E

Extreme Networks Special Edition

by David Coleman – CWNE #4

^{for}
dummies[®]
A Wiley Brand

Wi-Fi 6 & 6E For Dummies®, Extreme Networks Special Edition

Published by

John Wiley & Sons, Inc.

111 River St.

Hoboken, NJ 07030-5774

www.wiley.com

Copyright © 2022 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. Extreme Networks and the Extreme Networks logo are trademarks or registered trademarks of Extreme 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- 80787-2 (pbk); ISBN 978-1-119- 80788-9 (ebk)

Publisher's Acknowledgments

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

Development Editor: Ryan Williams

Project Manager: Jennifer Bingham

Acquisitions Editor: Ashley Coffey

Editorial Manager: Rev Mengle

Business Development

Representative: Molly Daugherty

Content Refinement Specialist:

Mohammed Zafar

Author Acknowledgements

Learning and working with Wi-Fi technology has been a long and adventurous journey for the past 22 years. Writing this book has been yet another exciting chapter of that journey, and I would like to thank a few people who have been Sherpas during this trek.

It has been an honor to join the visionaries at the Office of the CTO at Extreme Networks. Working with Nabil Bukhari (CTO), Markus Nispel, Alena Amir, Marcus Burton, Jon Filson, Tim Harrison, Ed Koehler, Yuri Ostrovsky, Kurt Semba, Alexey Reznik, Daniel Schieber, Kashif Siddiqui, Giacomo Bernardi, Divya Balu Pazhayannur, Doug McDonald, and the entire OCTO team. Many special thanks to my boss, Carla Guzzetti for bringing me into this elite group of professionals.

Certainly, I must thank all of my PLM and Wi-Fi co-workers at Extreme Networks, including Jeevan Patil, Rosalie Bibona, Andrew Garcia, Joe Zhao, Paulo Francisco, Mike Eversole, Gorden Jennings, Yarnin Israel, David van Schravendijk, Alexandra Gates, Deven Ducommun, Marko Tisler, Jeffrey Thorpe, Mike Lane, Gregor Vucajnk, Cristian Mercia, Wendy Kastner, Hardik Ajmera, Rich Hunt, and of course, Alan Cuellar Amrod.

As always, true heroes work in engineering! So many thanks to Kevin Lin, Vlado Tasev, Lucy Kestekian, Dionis Hristov, Hitendrasinh Gohil, David Zhang, Zhao Chen, Dmitry Shnayder, Akshay Ranjan, Roie Lev, Chao Zhu, Xindan Wang, TC Yang, Abhijit Chandavale, Mudit Seth, Arnab Ray, Robert Greenway, Ed Chee, Robert Koehler, Pow Yap, Gabriel Strasser, Koon Hung Lee, Phong Nguyen, Harish Vasista, Sangharsha Gangadharaswamy, Vijay Kumar Duba, Lagamanna Bawoor, Rajesh Pitchai, and many others.

Let's not forget a shout-out to Mark Dellavalle, Roy Verboeket, and James Sarikko for SE leadership. Your team in the field working with customers and partners is where the real Wi-Fi magic happens.

Special thanks to Kendra Lucciano, Nora Guzman, Blair Donald, Lisa Yeaton, and Christi Nicolacopoulos for all their support. Many thanks to my good friend Perry Correll.

I also need to give a huge shout-out to our friends at Broadcom, including Vijay Nagarajan, Mike Powell, Christopher Szymanski, Thomas Derham, Gabriel Desjardins, and Nitin Madan.

I must send out high-fives to both Claus Hetting of Wi-Fi NOW and Keith Parsons of WLANPros for the numerous contributions both organizations make to the Wi-Fi industry.

As always, thanks to my project editor, Jen Bingham, for her patience during the entire writing process.

Thank you to Tiago Rodrigues, Steve Namaseevayum, Bruno Thomas, and all of our friends at the Wireless Broadband Alliance.

Finally, many thanks to the entire staff at the Wi-Fi Alliance (www.wi-fi.org), especially Cari Eissler and Tina Hanzlik.

About the Author

David D. Coleman is the Director of Wireless Networking at the *Office of the CTO* for Extreme Networks. David is a technology evangelist, public speaker, and proficient author. David travels the world for both customer and channel partner engagements, speaking events and training sessions. He has instructed IT professionals from around the globe in Wi-Fi design, security, administration, and troubleshooting. David has written multiple books, blogs, and white papers about wireless networking, and he is considered an authority on 802.11 technology. Prior to working at Extreme, he specialized in corporate and government Wi-Fi training and consulting. In the past, he has provided Wi-Fi consulting for numerous private corporations, the US military, and other federal and state government agencies. David is also the 2020 recipient of the Wi-Fi Lifetime Achievement Award. When he is not traveling, David resides in Atlanta, Georgia and Chapala, Mexico. David is CWNE #4 and is the co-author of Sybex Publishing's Certified Wireless Network Administrator (CWNA) Study Guide – 6th Edition.

You can follow David on social media at:

Twitter: @mistermultipath

LinkedIn: <https://www.linkedin.com/in/mistermultipath/>

Introduction

In 2019, Wi-Fi technology celebrated its 20th birthday. It is hard to believe that Wi-Fi has been around for 22 years. And now, beginning in the year 2021, Wi-Fi is making a massive leap into the 6 GHz frequency band.

Previous generations of Wi-Fi focused on increasing data rates and speed. Wi-Fi 6 (also known as 802.11ax) is the new generation of Wi-Fi technology with a new focus on efficiency and performance. The Wi-Fi Alliance began certifying 802.11ax technology in August 2019, with a new certification called Wi-Fi CERTIFIED 6. Wi-Fi 6 technology is all about better and more efficient use of the existing radio frequency medium. Wi-Fi 6 handles client density more efficiently through a new channel-sharing capability that promises true multi-user communications on both the downlink and uplink. Wi-Fi 6 also uses a new client power-saving mechanism that schedules wake-times to improve client battery life.

In early 2020, the U.S. Federal Communications Commission (FCC) voted unanimously to make 1,200 megahertz of spectrum in the 6 GHz band available for unlicensed use in the United States. To put this in perspective, the new 6 GHz spectrum available for Wi-Fi is more than double the usable channels of the 2.4 GHz and 5 GHz channels combined. So effectively, it triples the available unlicensed spectrum available for Wi-Fi. This, my friends, is a big deal. In late 2020, the Wi-Fi Alliance announced Wi-Fi 6E as an “extension” for certifying the 802.11ax features and Wi-Fi 6 capabilities in the 6 GHz band. Wi-Fi 6E is the industry name that identifies Wi-Fi devices that operate in 6 GHz.

I am fond of saying that Wi-Fi technology is ingrained into our everyday lives. Wi-Fi has become an essential part of our daily worldwide communications culture. For over 20 years, Wi-Fi has provided true wireless mobility and secure connectivity in the enterprise. However, one of the lessons we have learned during the recent pandemic is that companies are changing the way they do business. For example, employees might no longer need to report to an office or campus for the entire five-day work week. Instead, they may prefer to work remotely, either part-time or full-time, but with expectations of the same enterprise Wi-Fi experience. Therefore, the enterprise no longer just resides in a building, but instead, the enterprise is wherever employees choose to work. We now have an *infinite enterprise*.

Both Wi-Fi 6 and Wi-Fi 6E are part of a wireless paradigm shift toward *infinitely distributed connectivity*. Enterprise companies need to connect anybody, anywhere, to any other person, device, or application. Wi-Fi 6 and 6E technology and the new 1,200 MHz of 6 GHz spectrum offer an enhanced user experience in all enterprise verticals, including K-12 and higher education, retail, manufacturing, and healthcare. On a personal note, I must say that I am very excited about the future of Wi-Fi because of the new availability of the 6 GHz spectrum and the technology of 802.11ax. Wi-Fi 6 and 6E moves us forward in the *infinite enterprise*. In this book, you learn about Wi-Fi 6 efficiency enhancements, the value of the technology as it moves into the 6 GHz spectrum, and the real-world implications of this historic evolution of Wi-Fi technology.

I think Rosalie Bibona, Senior WLAN Product Manager at Extreme Networks, put it best when she stated recently, “It’s almost like Wi-Fi is being born again!” I could not agree with her more, Wi-Fi 6E marks a new beginning for Wi-Fi, and the future is extremely bright.

About This Book

Wi-Fi 6 & 6E For Dummies, Extreme Networks Special Edition, consists of eight chapters that explore:

- » Wi-Fi Traffic Jam (Chapter 1)
- » The Secret Sauce of Wi-Fi 6: OFDMA (Chapter 2)
- » Holy Cow! - MU-MIMO (Chapter 3)
- » A Splash of Wi-Fi Color (Chapter 4)
- » Additional Wi-Fi 6 Enhancements (Chapter 5)
- » Wi-Fi 6E – A New Beginning in 6 GHz (Chapter 6)
- » Security in a 6 GHz Wi-Fi 6E World (Chapter 6)
- » Wi-Fi 6 and 6E Key Questions (Chapter 8)
- » Ten Things to Know about Wi-Fi 6 & 6E (Chapter 9)

Foolish Assumptions

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

Mainly, we assume that you are an IT infrastructure professional — someone with networking, wireless, or cloud in their title — and that you work for a medium to large organization or enterprise with robust Wi-Fi business requirements and you are interested in what is next for Wi-Fi.

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 Wi-Fi.

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 we sure hope you will appreciate these tips. This icon points out useful nuggets of information.

Beyond the Book

There is only so much I can cover in 96 short pages, so if you find yourself at the end of this book, thinking, “Where can I learn more?” A great reference guide about all things Wi-Fi is David Coleman and David Westcott’s vendor-neutral book, *CWNA*

Certified Wireless Network Administrator Study Guide: Exam CWNA-108 (Wiley). In the meantime, here are some links to some other great resources:

» **Extreme's Wi-Fi 6 web page:** <https://www.extremenetworks.com/WiFi6>.

» **David Coleman's Blog at Extreme Networks:** <https://www.extremenetworks.com/extreme-networks-blog/author/dcoleman>.

There are other great blog authors at Extreme Networks, so please check them out as well.

» **The Wi-Fi CERTIFIED 6™ resources from the Wi-Fi Alliance:**
<https://www.wi-fi.org/discover-wi-fi/wi-fi-certified-6>.

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 or backwards).

IN THIS CHAPTER

- » Looking at the timeline of Wi-Fi technology
- » Introducing 802.11ax and Wi-Fi 6
- » Considering the current Wi-Fi traffic jam
- » Recognizing that Wi-Fi 6 is about efficiency
- » Understanding the definition of multi-user
- » Extending Wi-Fi 6 into 6 GHz

Chapter 1

Wi-Fi Traffic Jam

In this chapter, you look back at the evolution of Wi-Fi standards and the challenges that exist in current Wi-Fi networks, as well as the newest generation of Wi-Fi technology that defines efficiency enhancements: Wi-Fi 6.

Wi-Fi Timeline

The Institute of Electrical and Electronics Engineers (IEEE) is the professional society that creates and maintains standards that we use for communications, such as the 802.3 Ethernet standard for wired networking. Since 1997, the IEEE has maintained the 802.11 standard for *wireless local area network (WLAN)* technology.



In your daily life, you probably instead use the familiar phrase *Wi-Fi* to discuss the same technology. Many people mistakenly assume that Wi-Fi is an acronym for the phrase *wireless fidelity* (much like hi-fi is short for *high fidelity*), but Wi-Fi is simply a brand name used to market 802.11 WLAN technology. The name that people will always recognize for the technology is Wi-Fi. As a matter of fact, Wi-Fi has become an essential part of our

daily worldwide communications culture. Wi-Fi technology is ingrained into our everyday lives.

The 802.11 working group has about 400 active members from more than 200 wireless companies. It consists of standing committees, study groups, and numerous *task groups*. Over the years, various 802.11 task groups have been in charge of revising and amending the original standard. Each group is assigned a letter from the alphabet, and it is common to hear the term 802.11 *alphabet soup* when referring to all the amendments created by the multiple 802.11 task groups. The goal of each amendment is to enhance 802.11 technology. In the past, many of the enhancements had a primary emphasis on higher data rates and faster speeds (see Table 1-1.) For example, 802.11b introduced data rates of up to 11 Mbps (megabits per second). 802.11a and 802.11g introduced data rates of 54 Mbps. 802.11n and 802.11ac enhanced data rates much further.

TABLE 1-1 **Wi-Fi Timeline**

Year	Amendment	Data rates	2.4 GHz	5 GHz	RF technology	Radios
1997	802.11 legacy	1 and 2 Mbps	✓		DSSS and FHSS	SISO
1999	802.11a	6 - 54 Mbps		✓	OFDM	SISO
1999	802.11b	1, 2, 5.5 and 11 Mbps	✓		HR-DSSS	SISO
2003	802.11g	6 - 54 Mbps	✓		OFDM	SISO
2009	802.11n	Up to 600 Mbps	✓	✓	OFDM	MIMO
2013	802.11ac	Up to 6.93 Gbps		✓	OFDM	MU-MIMO

In 1999, wireless was commercially introduced as a “nice to have” feature with the 802.11a and 802.11b ratifications. 802.11b, the most commonly used standard at the time, had very low speeds — only up to 11 Mbps (much lower than most Ethernet wired networks installed at the time) — but there were no Wi-Fi mobile devices and very few laptops, so 11 Mbps was fast enough. By 2003, Wi-Fi-enabled mobile devices were being introduced in

the market, and portable laptops became common for both business and personal use. The 802.11g standard was subsequently ratified, delivering up to 54 Mbps speeds on the 2.4 GHz frequency band.

In 2007, Apple introduced the first iPhone, and the smartphone became a modern reality. The 802.11n standard followed in 2009, delivering 100 Mbps of usable throughput. The 802.11n standard also brought about faster theoretical data rates of up to 600 Mbps and supported both 2.4 and 5 GHz devices. 802.11n was the last big paradigm shift in Wi-Fi technology when we switched from *single-input single-output (SISO)* radios to *multiple-input multiple-output (MIMO)* radios. We went from a time when an RF phenomenon known as multipath became constructive instead of destructive. By 2012, wireless mobile devices such as smartphones surpassed personal computer sales.

In 2013, 802.11ac expanded and simplified many of the technologies of 802.11n: Even higher data rates prevailed; however, 802.11ac only operates in the 5 GHz frequency band. Although data rates of up to 6.93 Gbps are theoretically possible with 802.11ac, in the real world, data rates of up to 400 to 600 Mbps are more likely. 802.11ac also introduced a multi-user technology known as multi-user MIMO (MU-MIMO); however, the implementation has been sparse.

As you can see, over the years, the main emphasis has been on faster speeds and higher data rates to meet the high-density demands in enterprise WLANs. However, there is a big misconception that data rates are the same as actual throughput. And furthermore, speed can be overrated. What good is a Ferrari that can travel at 300 km per hour if the Ferrari is stuck in traffic gridlock?

802.11ax = Wi-Fi 6

802.11ax-2020 is an IEEE amendment that defines modifications to the 802.11 Physical (PHY) layer and the Medium Access Control (MAC) sublayer for *high efficiency* operations in frequency bands between 1 GHz and 6 GHz. Much like very high throughput (VHT) is the technical term for 802.11ac, high efficiency (HE) is the technical term for 802.11ax.

The Wi-Fi Alliance (www.wi-fi.org) is a global, nonprofit industry association of about 550 member companies devoted to promoting the growth of WLANs. One of the primary tasks of the Wi-Fi Alliance is to market the Wi-Fi brand and raise consumer awareness of new 802.11 technologies as they become available. The Wi-Fi Alliance's main task is to ensure the interoperability of WLAN products by providing certification testing. Products that pass the Wi-Fi certification process receive a *Wi-Fi Interoperability Certificate* that includes detailed information about the individual product's Wi-Fi certifications.



REMEMBER

Recently, the Wi-Fi Alliance adopted a new generational naming convention for Wi-Fi technologies. The goal is that the new naming convention will be easier to understand for the average consumer as opposed to the alphabet-soup naming used by the IEEE. The Wi-Fi Alliance began certifying 802.11ax technology in August 2019 with a new certification called Wi-Fi CERTIFIED 6. Because 802.11ax technology is such a major paradigm shift from previous versions of 802.11 technology, it has been bestowed with the generational name of *Wi-Fi 6*. Older versions of 802.11 technology also align with this new naming convention. For example, 802.11ac can be referenced as Wi-Fi 5, and 802.11n is Wi-Fi 4, as shown in Figure 1-1.




FIGURE 1-1: Generations of Wi-Fi.



TIP

Throughout this book, I use both the IEEE 802.11 terminology and the Wi-Fi Alliance generational terminology. 802.11ax and Wi-Fi 6 mean the same thing, but the term Wi-Fi 6 will be more prevalent with the general population. Geeky WLAN professionals might use term 802.11ax, while your grandma will understand the generational name of Wi-Fi 6.

Wi-Fi Traffic Congestion

Although Wi-Fi is a resilient technology, it has not necessarily been efficient. Wi-Fi operates at both layer 1 and layer 2 of the OSI model and the inefficiency exists at both layers.



REMEMBER

Historically, previous 802.11 amendments defined technologies that gave us higher data rates and wider channels but did not address efficiency. An often-used analogy is that faster cars and bigger highways have been built, but traffic jams still exist. Despite the higher data rates and 40/80/160 MHz channels used by 802.11n/ac radios, multiple factors contribute to the Wi-Fi traffic congestion, which do not provide for an efficient use of the medium.

So why exactly is there a Wi-Fi traffic jam? 802.11 data rates are not TCP throughput. Always remember that *radio frequency (RF)* is a half-duplex medium and that the 802.11 medium contention protocol of CSMA/CA consumes much of the available bandwidth. In laboratory conditions, TCP throughput of 60 to 70 percent of the operational data rate can be achieved using 802.11n/ac communication between one access point (AP) and one client. The aggregate throughput numbers are considerably less in real-world environments with the active participation of multiple Wi-Fi clients communicating through an AP. As more clients contend for the medium, the medium contention overhead increases significantly, and efficiency drops. Therefore, the aggregate throughput is usually at best 50 percent of the advertised 802.11 data rate. Not very efficient.

What else contributes to Wi-Fi traffic congestion? Because legacy Wi-Fi clients often still participate in enterprise, RTS/CTS protection mechanisms are needed. This contributes to the inefficiency. Figure 1-2 shows that about 60 percent of all Wi-Fi traffic is 802.11 control frames, and 15 percent is 802.11 management frames. Control and management frames consume 75 percent of the usable airtime, and only 25 percent of Wi-Fi traffic is used for 802.11 data frames. Additionally, layer 2 retransmissions resulting from either RF interference or a poorly designed WLAN, can also contribute to Wi-Fi inefficiency.




FIGURE 1-2: 802.11 traffic.

High data rates are useful for a large data payload; however, the bulk of 802.11 data frames (75–80 percent) are small and under 256 bytes, as shown in Figure 1-3. Each small frame requires a PHY header, a MAC header, and a trailer. The result is excessive PHY/MAC overhead as well as medium contention overhead for each small frame. Small frames can be aggregated to reduce the overhead; however, in most cases, the small frames are not aggregated, because they must be delivered sequentially due to the higher layer application protocols. For example, VoIP packets cannot be aggregated, because they must arrive sequentially.

Despite the higher data rates and wide channels that can be used by 802.11n/ac radios, the result is Wi-Fi traffic congestion. Automobile traffic congestion can result in drivers becoming frustrated and thereby engaging in road rage. Wi-Fi 6 (802.11ax) technology is all about better 802.11 traffic management and hopefully eliminating Wi-Fi radio rage.




FIGURE 1-3: 802.11 data frame size.

Wi-Fi 6 — More than Just Speeds and Feeds

Wi-Fi 6 (802.11ax) technology is all about better and more efficient use of the existing radio frequency medium. Higher data rates are not the primary goal of Wi-Fi 6. The goal is better and more efficient 802.11 traffic management. Figure 1-4 depicts many of the new enhancements defined for Wi-Fi CERTIFIED 6 radios. Most of these Wi-Fi 6 capabilities will be discussed in great detail throughout this book.

Most of the Wi-Fi 6 enhancements are at the PHY layer and involve a new multi-user version of *orthogonal frequency-division multiplexing (OFDM)* technology instead of the single-user OFDM technology already used by 802.11a/g/n/ac radios.

Another significant Wi-Fi 6 change is that an AP can actually supervise both downlink and uplink transmissions to multiple client radios while the AP has control of the medium. Beamforming permits the AP greater precision in sending signals to specific devices while preventing interference from others. BSS coloring

provides an identifier that helps devices differentiate its traffic from data associated with other devices.

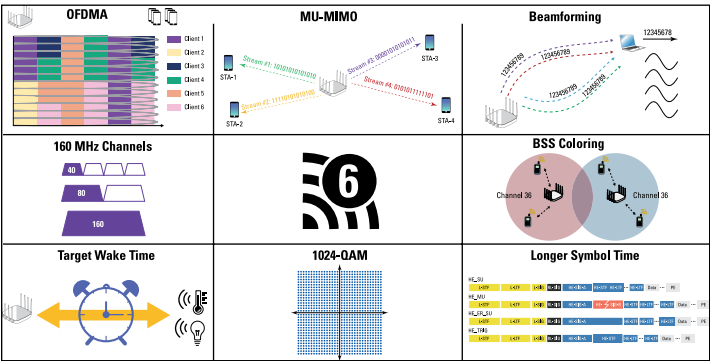


FIGURE 1-4: Wi-Fi 6 capabilities overview.

In addition to these multi-user efficiency enhancements, Wi-Fi 6 (802.11ax) radios will be backward compatible with 802.11/a/b/g/n/ac radios that operate in the 2.4 and 5 GHz frequency band. Backward compatibility does not exist in the 6 GHz frequency band and is not needed. Table 1-2 shows a high-level comparison of 802.11n, 802.11ac, and 802.11ax capabilities. Please note that unlike 802.11ac radios, which can transmit only on the 5 GHz frequency band, 802.11ax radios can operate on the 2.4 GHz and 5 GHz frequency bands and the new 6 GHz frequency band.

TABLE 1-2 802.11n, 802.11ac, and 802.11ax comparison

	802.11n (Wi-Fi 4)	802.11ac (Wi-Fi 5)	802.11ax (Wi-Fi 6)
Frequency bands	2.4 GHz and 5 GHz	5 GHz only	2.4 GHz, 5 GHz, 6 GHz
Channel size (MHz)	20, 40	20, 40, 80, 80 + 80, and 160	20, 40, 80, 80 + 80, and 160
Frequency multiplexing	OFDM	OFDM	OFDM and OFDMA
Subcarrier spacing (KHz)	312.5	312.5	78.125
OFDM symbol time (μs)	3.2	3.2	12.8
Guard interval (μs)	.04 or .08	.04 or .08	.08, 1.6, or 3.2

	802.11n (Wi-Fi 4)	802.11ac (Wi-Fi 5)	802.11ax (Wi-Fi 6)
Total symbol time (μs)	3.6 or 4.0	3.6 or 4.0	13.6, 14.4, or 16.0
Modulation	BPSK, QPSK, 16-QAM, 64-QAM	BPSK, QPSK, 16-QAM, 64-QAM, 256-QAM	BPSK, QPSK, 16-QAM, 64-QAM, 256-QAM, 1024-QAM
MU-MIMO	N/A	DL	DL and UL
OFDMA	N/A	N/A	DL and UL

As you can see in Table 1-2, Wi-Fi 6 does support 40 MHz, 80 MHz, and 160 MHz channels. For the bulk of the discussion of Wi-Fi 6 communication, however, I will focus on 20 MHz channels. That being said, a discussion about 40 MHz, 80 MHz, and 160 MHz channels in the 6 GHz frequency band will be an essential topic in this book.

Multi-User

The term *multi-user (MU)* simply means that transmissions between an AP and multiple clients can occur at the same time, depending on the supported technology. However, the MU terminology can be very confusing when discussing Wi-Fi 6. MU capabilities exist for both OFDMA and MU-MIMO. I will explain the key differences later in this book.

Wi-Fi 6 makes use of both multi-user technologies, OFDMA and MU-MIMO. Do not confuse OFDMA with MU-MIMO. OFDMA allows for multiple-user access by subdividing a channel (see Chapter 2). MU-MIMO allows for multiple-user access by modulating different spatial streams of data (see Chapter 3). If I reference the car and road analogy discussed earlier, OFDMA uses a single road subdivided into multiple lanes for use by different cars at the same time, whereas MU-MIMO uses different single lane roads to arrive at the same destination.



When discussing Wi-Fi 6, there is often a lot of confusion because many people may already be somewhat familiar with MU-MIMO technology introduced with 802.11ac (Wi-Fi 5). What most people are not familiar with is the multi-user technology of OFDMA. Most of the efficiency benefits of Wi-Fi 6 are a result of multi-user OFDMA.

Extending Wi-Fi 6 into 6 GHz

In early 2020, the U.S. Federal Communications Commission (FCC) voted unanimously to make 1,200 megahertz of spectrum in the 6 GHz band available for unlicensed use in the United States. And the great news is that we can now use 6 GHz for Wi-Fi. To put this in perspective, the new 6 GHz spectrum available for Wi-Fi is more than double the usable channels of the 2.4 GHz and 5 GHz channels combined. So effectively, it triples the available unlicensed spectrum available for Wi-Fi.

Just look at all the new Wi-Fi channels in 6 GHz, as shown in Figure 1-5. That is a truckload of new channels! In the United States, there are as many as 59 new 20 MHz channels available across four U-NII bands. Additionally, there is the potential to use 29 new 40 MHz channels. Even more surprising is that the 14 new 80 MHz channels in the 6 GHz band are expected to be used extensively in the enterprise. Once again, if I reference the car and road analogy, 6 GHz is a superhighway.




FIGURE 1-5: 6 GHz frequency band.

The great news is that many other world regions are also working toward making all or portions of the 6 GHz frequency band available for Wi-Fi. As of today, over 50 countries have also adopted new regulations for the unlicensed use of 6 GHz. Ratification for unlicensed use usually follows regulatory approval in a timely fashion.

In late 2020, the Wi-Fi Alliance announced Wi-Fi 6E as an “extension” for certifying the 802.11ax features and capabilities of Wi-Fi 6 to the 6 GHz band. Wi-Fi 6E is the industry marketing name that identifies Wi-Fi devices that operate in 6 GHz.

As shown in Figure 1-6, Wi-Fi has an estimated global economic impact of \$3.3 trillion dollars in 2021. And the potential of 1,200 MHz of new frequency space for Wi-Fi communications is mind-boggling. Opening the 6 GHz frequency for Wi-Fi communication is expected to bring \$154 billion in economic value to the United States by 2025.

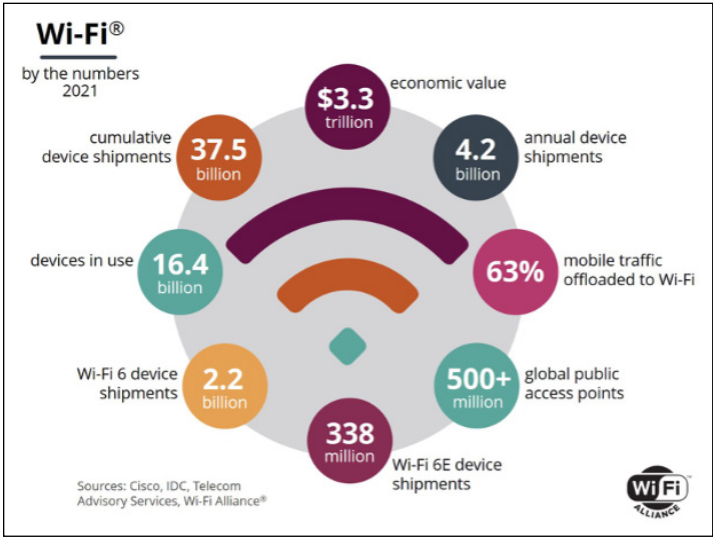


FIGURE 1-6: Wi-Fi by the numbers in 2021.

The availability of the 6 GHz frequency space for Wi-Fi will have all kinds of real-world implications. The 2.4 GHz band is considered a best-effort band, and now 5 GHz can also be crowded. 6 GHz will enable crystal-clear communications because it has

never been used before for Wi-Fi. I see immediate value for connecting mission-critical IoT devices, such as healthcare patient monitoring equipment, in this new *pristine* Wi-Fi environment. One of the other immediate use cases expected for 6 GHz Wi-Fi will be indoor mesh backhaul. For access, legacy 2.4 GHz and 5 GHz clients will connect to the same frequency radios in the Wi-Fi 6 APs. If needed, the 6 GHz radios in the APs can transport the legacy client traffic via an indoor 6 GHz mesh backhaul link.

In the long term, I predict tremendous development and innovation for higher bandwidth applications. We are at the beginning of a renaissance of innovation for virtual reality and augmented reality applications that can be used via a Wi-Fi connection. For example, an employee may use an augmented reality application to emulate an enterprise office experience while working remotely. As a direct result of Wi-Fi 6E, we are poised to see new methods of user engagement that we previously have only seen in science fiction movies.

Even though Wi-Fi 6 is now two years mature, porting 802.11ax technology to the 6 GHz frequency band will not be without its challenges. Well-defined indoor and outdoor power regulations are being established to protect existing 6 GHz incumbents. Because there are so many 6 GHz channels, new AP discovery mechanisms exist, so Wi-Fi 6E clients can connect and roam to Wi-Fi 6E APs. Much thought needs to be given to the 6 GHz WLAN design, including range, channel reuse, channel size, and primary channel selection. New Wi-Fi security mandates for 6 GHz will also have implications. Because there is so much to learn, I have written an entire chapter dedicated to all these Wi-Fi 6E considerations.

Wi-Fi 6 brings us the needed efficiency enhancements to ease the existing Wi-Fi traffic jam. And Wi-Fi 6E gives us a brand new 6 GHz superhighway. Wi-Fi 6 has been an essential technology update, and now Wi-Fi 6E is a monumental spectrum update.

- » Differentiating between OFDM and OFDMA
- » Understanding resource units, trigger frames, and buffer status reports
- » Recognizing the efficiency benefits of uplink and downlink OFDMA
- » Explaining operating mode indication

Chapter 2

The Secret Sauce of Wi-Fi 6: OFDMA

If you can serve more than one person (or device) at a time, why not keep everybody happy and do just that? Older Wi-Fi technology may be restricted to serving a single user per channel, but this is the future. Let's open the channels up! In this chapter, you learn about a multi-user technology called OFDMA, which is the main ingredient of Wi-Fi 6 technology.

OFDMA

Orthogonal frequency division multiple access (OFDMA) is arguably the most important new Wi-Fi 6 capability. Legacy 802.11a/g/n/ac radios currently use orthogonal frequency division multiplexing (OFDM) for single-user transmissions on any given channel. Wi-Fi 6 radios utilize orthogonal frequency division multiple access (OFDMA), which is a multi-user version of the OFDM digital-modulation technology. OFDMA subdivides a Wi-Fi channel into smaller frequency allocations, called *resource units (RUs)*, thereby enabling an access point (AP) to synchronize communication (uplink and downlink) with multiple individual clients

assigned to specific RUs. By subdividing the channel, small frames can be simultaneously transmitted to multiple users in parallel.



TIP

OFDMA is ideal for most network applications and results in better frequency reuse, reduced latency, and increased efficiency.

Think of OFDMA as a technology that partitions a Wi-Fi channel into smaller subchannels so that simultaneous multiple-user transmissions can occur. For example, a traditional 20 MHz channel might be partitioned into as many as nine smaller subchannels. Using OFDMA, a Wi-Fi 6 AP could simultaneously transmit small frames to nine Wi-Fi 6 clients. The Wi-Fi CERTIFIED 6 certification program from the Wi-Fi Alliance currently validates the simultaneous transmission of up to four resource units. Wouldn't you prefer four checkout aisles in the supermarket versus a single line?

OFDMA is a much more efficient use of the medium for smaller frames. The simultaneous transmission cuts down on excessive overhead at the medium access control (MAC) sublayer, as well as medium contention overhead. The AP can allocate the whole channel to a single user or partition it to serve multiple users simultaneously, based on client traffic needs. The goal of OFDMA is better use of the available frequency space. OFDMA technology has been time-tested with other RF communications. For example, OFDMA is used for downlink LTE cellular communications.

Subcarriers

Both OFDM and OFDMA divide a channel into subcarriers through a mathematical function known as an *inverse fast Fourier transform (IFFT)*. The spacing of the subcarriers is orthogonal, so they do not interfere with one another despite the lack of guard bands between them. This creates signal nulls in the adjacent subcarrier frequencies, thus preventing *intercarrier interference (ICI)*.

What are some of the key differences between OFDM and OFDMA? As shown in Figure 2-1, a 20 MHz 802.11n/ac channel consists of 64 subcarriers. Fifty-two of the subcarriers are used to carry modulated data; four of the subcarriers function as pilot carriers; and eight of the subcarriers serve as guard bands. OFDM subcarriers are sometimes also referred to as OFDM *tones*. In this book, I use both terms interchangeably. Each OFDM subcarrier is 312.5 KHz.




FIGURE 2-1: 802.11n/ac 20 MHz channel – OFDM subcarriers.

802.11ax introduces a longer OFDM symbol time of 12.8 microseconds, which is four times the legacy symbol time of 3.2 microseconds (for more information on symbol time, see Chapter 5). As a result of the longer symbol time, the subcarrier size and spacing decreases from 312.5 kHz to 78.125 kHz. The narrower subcarrier spacing allows better equalization and enhanced channel robustness. Figure 2-2 shows the difference between OFDM and OFDMA subcarrier size and spacing. Because of the 78.125 kHz spacing, an OFDMA 20 MHz channel consists of a total of 256 subcarriers (tones). Which tones function to carry modulated data and which




FIGURE 2-2: OFDM and OFDMA subcarrier spacing.

tones serve as pilot carriers is dependent on the size and location of allocated resource units (RUs).



TECHNICAL
STUFF

Just like with OFDM, there are three types of subcarriers for OFDMA, as follows:

- » **Data subcarriers:** These subcarriers will use the same modulation and coding schemes (MCSs) as 802.11ac as well as two new MCSs with the addition of 1024 quadrature amplitude modulation (1024-QAM).
- » **Pilot subcarriers:** The pilot subcarriers do not carry any modulated data; however, they are used for synchronization purposes between the transmitter and receiver.
- » **Unused subcarriers:** The remaining unused subcarriers are mainly used as guard carriers or null subcarriers against interference from adjacent channels or subchannels.

With OFDMA, these tones are grouped together into partitioned subchannels, known as resource units (RUs). By subdividing the channel, parallel transmissions of small frames to multiple users can happen simultaneously. The data and pilot subcarriers within each resource unit are both adjacent and contiguous within an OFDMA channel.



TIP

For backward compatibility, Wi-Fi 6 radios still support OFDM. Keep in mind that 802.11 management and control frames will still be transmitted at a basic data rate using OFDM technology that legacy 802.11a/g/n/ac radios can understand. Therefore, management and control frames are transmitted across all the OFDM subcarriers of an entire primary 20 MHz channel. OFDMA is only for 802.11 data frame exchanges between Wi-Fi 6 APs and Wi-Fi 6 clients.

Resource Units (RUs)

When an 802.11n/ac AP transmits downstream to 802.11n/ac clients on an OFDM channel, the entire frequency space of the channel is used for each independent downlink transmission. In the example shown in Figure 2-3, the AP transmits to six clients independently over time. When using a 20 MHz OFDM channel, all of the 64 subcarriers are used for each independent transmission.

In other words, the entire 20 MHz channel is needed for the communication between the AP and a single OFDM client. The communications are *single-user*. The same holds true for any uplink transmission from a single 802.11n/ac client to the 802.11n/ac AP. The entire 20 MHz OFDM channel is needed for the client transmission to the AP.




FIGURE 2-3: OFDM transmissions over time.

As previously stated, an OFDMA channel consists of a total of 256 subcarriers (tones). These tones can be grouped into smaller subchannels known as *resource units (RUs)*. As shown in Figure 2-4, when subdividing a 20 MHz channel, a Wi-Fi 6 AP can designate 26, 52, 106, and 242 subcarrier resource units (RUs), which equates roughly to 2 MHz, 4 MHz, 8 MHz, and 20 MHz channels, respectively. The Wi-Fi 6 AP dictates how many RUs are used within a 20 MHz channel and different combinations can be used. The AP may allocate the whole channel to only one client at a time or it may partition the channel to serve multiple clients simultaneously. For example: A Wi-Fi 6 AP could simultaneously communicate with one Wi-Fi 6 client using 8 MHz of frequency space, while communicating with three other Wi-Fi 6 clients using 4 MHz subchannels. These simultaneous communications can be either downlink or uplink.

In the example shown in Figure 2-5, the Wi-Fi 6 AP first simultaneously transmits downlink to Wi-Fi 6 clients 1 and 2. The 20 MHz OFDMA channel is effectively partitioned into two subchannels. Remember that an OFDMA 20 MHz channel has a total



FIGURE 2-4: OFDM resource units – 20 MHz channel.

of 256 subcarriers; however, the AP is simultaneously transmitted to clients 1 and 2 using two different 106-tone resource units. In the second transmission, the AP simultaneously transmits downlink to clients 3, 4, 5, and 6. In this case, the OFDMA channel had to be partitioned into four different 52-tone resource units. In the third transmission, the AP uses a single 242-tone resource unit to transmit downlink to a single client (client 5). Using a single 242-tone resource unit is effectively using the entire 20 MHz channel. In the fourth transmission, the AP simultaneously transmits downlink to clients 4 and 6 using two 106-tone resource units. In the fifth transmission, the AP once again only transmits downlink to a single client with a single RU utilizing the entire 20 MHz channel. In the sixth transmission, the AP simultaneously transmits downlink to clients 3, 4, and 6. In this instance, the 20 MHz channel is partitioned into three subchannels. Two 52 RUs are used for clients 3 and 4, and a 106-tone RU is used for client 6.



TECHNICAL
STUFF

The rules of medium contention still apply, though. The AP still must compete against legacy clients for a *transmission opportunity (TXOP)*. Once the AP has a TXOP, the AP is then in control of multiple Wi-Fi 6 clients for either downlink or uplink transmissions. The number of RUs used can vary on a per-TXOP basis. The AP assigns RUs to associated clients on a per-TXOP basis to maximize the download and upload efficiency.

As shown in Figure 2-5, the AP can partition the 20 MHz OFDMA channel on a continuous basis for downlink transmissions. Current-generation Wi-Fi 6 radios can subdivide a 20 MHz channel into as many as four 52-tone resource units. Resource units




FIGURE 2-5: OFDMA transmissions over time.

can also be used to subdivide 40 MHz or 80 MHz channels. An AP can also partition 40 MHz, 80 MHz, and even 160 MHz channels into various combinations of RUs, as shown in Table 2-1. For example, if an 80 MHz channel was subdivided using strictly 26 subcarrier RUs, 37 Wi-Fi 6 clients could theoretically communicate simultaneously using their OFDMA capabilities. Quite frankly, it's doubtful that even large channels will be partitioned among more than four clients. However, resource unit allocation and scheduling mechanisms are likely to become more precise in later generations of chipsets and firmware. Remember, the whole point of OFDMA is to make use of smaller subchannels.

TABLE 2-1 Resource Units and Wide Channels

Resource Units (RUs)	20 MHz channel	40 MHz channel	80 MHz channel	160 MHz channel	80 + 80 MHz channel
996 (2x) subcarriers	n/a	n/a	n/a	1 client	1 client
996 subcarriers	n/a	n/a	1 client	2 clients	2 clients
484 subcarriers	n/a	1 client	2 clients	4 clients	4 clients
242 subcarriers	1 client	2 clients	4 clients	8 clients	8 clients
106 subcarriers	2 clients	4 clients	8 clients	16 clients	16 clients
52 subcarriers	4 clients	8 clients	16 clients	32 clients	32 clients
26 subcarriers	9 clients	18 clients	37 clients	74 clients	74 clients



The Wi-Fi CERTIFIED 6 certification program from the Wi-Fi Alliance currently validates up to four resource units. Initially, most real-world Wi-Fi 6 deployments use 20 MHz or 40 MHz channels with a maximum of four clients participating in multi-user OFDMA transmissions per TXOP. For example, a 20 MHz channel is subdivided into four 52-tone resource units, or a 40 MHz channel is partitioned into four 106-tone resource units. It's likely that subdividing 80 MHz and even 160 MHz channels using OFDMA will become more prevalent with the advent of the 1,200 MHz frequency space as Wi-Fi 6 (802.11ax) technology moves into the 6 GHz band.

Trigger Frames

When referencing downlink and uplink OFDMA transmissions, you see the acronyms of DL-OFDMA and UL-OFDMA pop up (because we just don't have enough acronyms). In both cases, *trigger frames* bring about the necessary frame exchanges for multi-user communications. For example, an AP uses trigger frames to allocate OFDMA RUs to Wi-Fi 6 clients. Multiple types of 802.11 control frames can function as trigger frames including basic trigger frames, multi-user request-to-send (MU-RTS) frames, buffer status report (BSRP) frames, and more.



RU allocation information is communicated to clients at both the PHY and MAC layers. At the Physical layer, RU allocation information can be found in the HE-SIG-B field of the PHY header of a trigger frame. The HE-SIG-B field is used to communicate RU assignments to clients. As shown in Figure 2-6, the HE-SIG-B field consists of two subfields: the common field and user-specific field. A subfield of the common field is used to indicate how a channel is partitioned into various RUs. For example, a 20 MHz channel might be subdivided into one 106-tone RU and four 26-tone RUs. The user-specific field comprises multiple-user fields that are used to communicate which users are assigned to each individual RU.

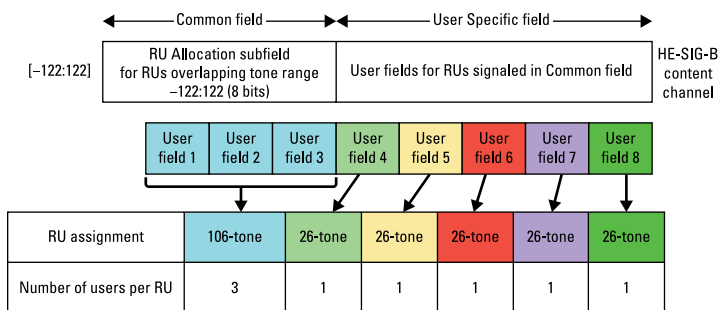


FIGURE 2-6: RU allocation at the PHY layer.

What about how RU allocation information is communicated at the MAC layer? RU allocation information is delivered in the *user information* field in the body of a trigger frame. Figure 2-7 displays a table of how RU allocation information is communicated at the MAC layer. The table highlights all the possible RUs within a 20 MHz channel and the subcarrier range for each RU. Each specific RU is defined by a unique combination of 7 bits within the user information field of the trigger frame, known as the RU allocation bits.

26 tone RU	RU-1	RU-2	RU-3	RU-4	RU-5	RU-6	RU-7	RU-8	RU-9
Subcarrier range	-121:-96	-95:-70	-68:-43	-42:-17	-16:-4, 4:16	17:42	43:68	70:95	96:121
RU allocation bits	0000000	0000001	0000010	0000011	0000100	0000101	0000110	0000111	0001000
52 tone RU	RU-1		RU-2			RU-3		RU-4	
Subcarrier range	-121:-70		-68:-17			17:68		70:121	
RU allocation bits	0100101		0100110			0100111		0101000	
106 tone RU	RU-1					RU-2			
Subcarrier range	-122:-17					17:122			
RU allocation bits	0110101					0110110			
242 tone RU	RU-1								
Subcarrier range	-122:-2, 2:122								
RU allocation bits	0111101								

FIGURE 2-7: RU allocation at the MAC layer.

In the example in Figure 2-8, the trigger frame allocates specific RUs to three client stations for simultaneous uplink transmission within a 20 MHz OFDMA channel. Clients STA-1 and STA-2 are each assigned to a 52-tone RU, whereas client STA-3 is assigned to a 106-tone RU.

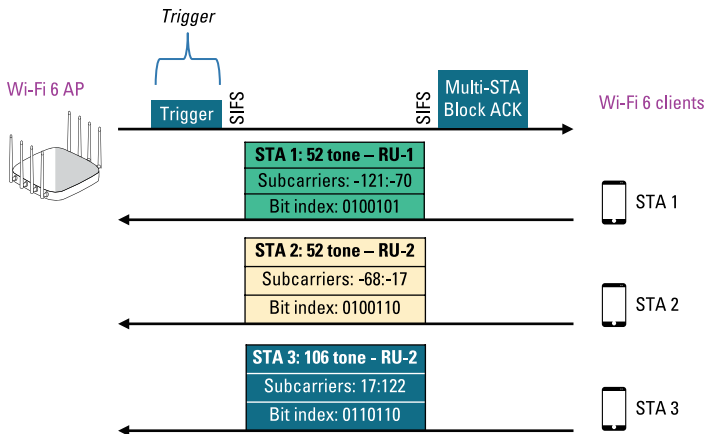


FIGURE 2-8: RU allocation with trigger frames.



For UL-OFDMA, a trigger frame sent by the AP tells the clients how many spatial streams and which modulation and coding scheme (MCS) to use when transmitting uplink on their assigned RUs. APs can also use trigger frames to tell clients to adjust their power settings for synchronized uplink transmissions. Please note that a Wi-Fi 6 client might be unable to satisfy the target RSSI due to its hardware or regulatory limitations.

DL-OFDMA

Take a look at how multi-user DL-OFDMA communications between a Wi-Fi 6 AP and Wi-Fi 6 clients.



OFDMA is only for 802.11 data frame exchanges between 802.11ax APs and 802.11ax clients. OFDMA is not used for management or control frames.

A Wi-Fi 6 AP will first need to contend for the medium and win a TXOP for the entire DL-OFDMA frame exchange. As shown in Figure 2-9, once an AP has won a TXOP, the AP might send a *multi-user request-to-send (MU-RTS)* frame. The MU-RTS frame has two purposes:

- » **Reserve the medium:** The MU-RTS frame is transmitted using OFDM (not OFDMA) across the entire 20 MHz channel so that legacy clients can also understand the MU-RTS. The

duration value of the MU-RTS frame is needed to reserve the medium and reset the NAV timers of all legacy clients for the remainder of the DL-OFDMA frame exchange. The legacy clients must remain idle while the multi-user OFDMA data frames are transmitted between the Wi-Fi 6 AP and the Wi-Fi 6 clients.

» **RU allocation:** The MU-RTS frame is also an extended trigger frame from the AP used to synchronize uplink clear-to-send (CTS) client responses for Wi-Fi 6 clients. The AP uses the MU-RTS as a trigger frame to allocate RUs. The Wi-Fi 6 clients send CTS responses in parallel using their assigned RUs.

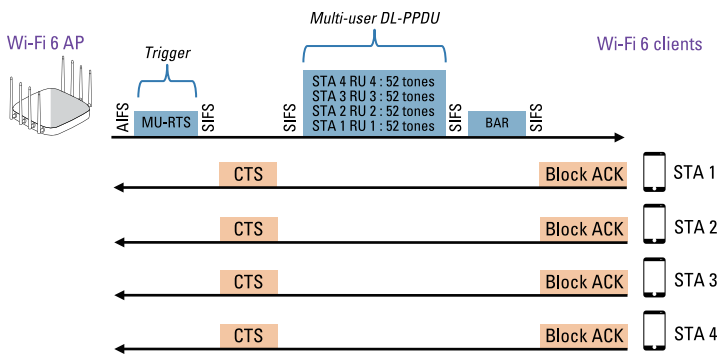


FIGURE 2-9: Downlink OFDMA.

After the parallel CTS response from the clients, the AP begins multi-user DL-PPDU transmissions from the AP to the OFDMA-capable clients. Keep in mind that the AP determined how to partition the 20 MHz channel into multiple RUs. Once the Wi-Fi 6 clients receive their data via their assigned RUs, they need to send a Block ACK to the AP. The AP will send a Block ACK request (BAR) frame followed by the clients replying with Block ACKs in parallel. Optionally, an automatic Block ACK can be sent by the clients in parallel.

Once the frame exchange is over, the AP or clients that win the next TXOP will then be able to transmit on the medium. Single-user communications can still occur for legacy clients. For example, if an 802.11n/ac client wins the next TXOP, the 802.11n/ac client will use an entire 20 MHz channel for an uplink transmission to the AP using OFDM.

UL-OFDMA

In the original 802.11 standard, the IEEE proposed an operational mode called *Point Coordination Function (PCF)*, which defined operations where the AP could control the medium for uplink client transmission. With PCF mode, the AP could poll clients for uplink transmissions during a contention-free period of time when the AP controlled the medium. However, PCF never caught on and was never implemented in the real world. 802.11ax now introduces mechanisms where the AP can once again control the medium for uplink transmissions using UL-OFDMA. You should understand that UL-OFDMA has nothing to do with PCF; the methods are very different. You should also understand that the 802.11ax AP must first contend for the medium and win a TXOP. Once the 802.11ax AP wins a TXOP, it can then coordinate uplink transmissions from 802.11ax clients that support UL-OFDMA.

UL-OFDMA is more complex than DL-OFDMA and may require the use of as many as three trigger frames. Each trigger frame is used to solicit a specific type of response from the Wi-Fi 6 clients. UL-OFDMA also requires the use of *buffer status report (BSR)* frames from the clients. Clients use BSR frames to inform the AP about the client's buffered data and about the QoS category of data. The information contained in BSR frames assists the AP in allocating RUs for synchronized uplink transmissions. The AP uses the information gathered from the clients to build uplink window times, client RU allocation, and client power settings for each RU. BSRs can be unsolicited or solicited. If solicited, the AP polls the clients for BSRs.

Now, look at how multi-user UL-OFDMA communications can occur between an AP and the clients. The Wi-Fi 6 AP will first need to contend for the medium and win a TXOP for the entire UL-OFDMA frame exchange. As shown in Figure 2-10, once a Wi-Fi 6 AP has won a TXOP, the AP will send the first trigger frame. A *buffer status report poll (BSRP)* frame is used to solicit information from the Wi-Fi 6 clients about their need to send uplink data. The clients will then respond with BSRs. The whole purpose of the BSR information is so the Wi-Fi 6 clients can assist a Wi-Fi 6 AP to allocate uplink multi-user resources. The AP will use this information to decide how to best allocate RUs to the clients for synchronized uplink transmissions.

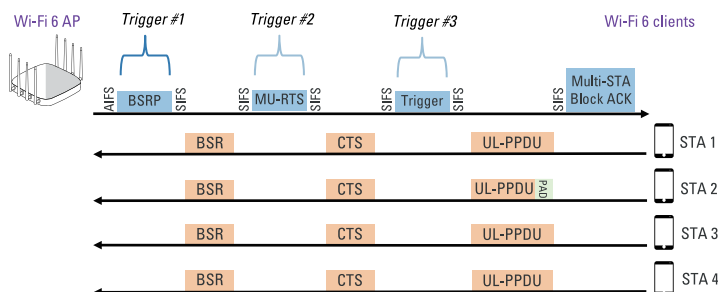


FIGURE 2-10: Uplink OFDMA.

If legacy clients exist, the AP may send a multi-user request-to-send (MU-RTS) frame, which functions as a second type of trigger frame. The RTS/CTS process is once again used to reserve the medium for OFDMA communications only.

A third and final *basic trigger frame* is needed to signal the Wi-Fi 6 clients to begin uplink transmission of their data with their assigned RUs. The basic trigger frame also dictates the length of the uplink window. The uplink client devices must all start and stop at the same time. The basic trigger frame also contains power control information so that individual clients can increase or decrease their transmit power. This data helps equalize the received power to the AP from all uplink clients and improve reception. Once the uplink data is received from the clients, the AP sends a single *multi-user Block ACK* to the clients. The AP also has the option of sending separate Block ACKs to each individual client.



TIP

All three trigger frames may or may not be needed for uplink transmission. For example, the MU-RTS trigger frame is only needed for protection mechanism purposes for legacy clients. The 6 GHz frequency band for 802.11ax technology does not require backward compatibility. Because 802.11a/b/g/n/ac radios operate on either the 2.4 GHz or the 5 GHz band, and not the 6 GHz band, there is no need for RTS/CTS protection mechanisms. Think of an express train or an HOV traffic lane. The 6 GHz frequency band will be a “pure” 802.11ax band for Wi-Fi 6E communication.

Buffer Status Reports

Wi-Fi 6 APs require specifics on client buffer states to perform appropriate synchronized uplink scheduling. As a result, the Wi-Fi 6 clients deliver BSRs to assist the AP in allocating uplink multi-user resources. Clients have two methods of delivering their buffer state information to the AP. Clients can *explicitly* deliver BSRs to the AP in response to a BSRP trigger frame (solicited BSR), as shown in Figure 2-10. However this solicited polling process does generate overhead. To minimize overhead, an AP can include a BSRP trigger frame together with other control, data, and management frames in one A-MPDU sent to a client that supports the capability.

Clients can *implicitly* deliver BSRs in the QoS Control field or BSR Control field of any frame transmitted to the AP (unsolicited BSR). Wi-Fi 6 clients can report unsolicited buffer status information for any given QoS class of traffic in any QoS Data or QoS Null frames it transmits. Additionally, as depicted in Figure 2-11, a Wi-Fi 6 client can report the buffer status of multiple QoS access categories using A-MPDU frame aggregation of QoS Data or QoS Null frames. The unsolicited BSR process is more efficient because it eliminates the need for polling.




FIGURE 2-11: Unsolicited buffer status reports and A-MPDU.

OMI

For backward compatibility purposes, legacy 802.11a/b/g/n/ac Wi-Fi clients still contend for the medium and win their own TXOP if they want to transmit uplink. This process may sound like aging heavyweight fighters in a title bout, but be assured it's quicker and way less violent. However, the uplink transmissions of Wi-Fi 6 clients are synchronized and controlled by the AP. I'm

often asked, “Can a Wi-Fi 6 client station suspend participation for synchronized uplink OFDMA and contend for the medium for an independent uplink transmission?” And I answer “Yes!” and launch directly into this elaboration.

802.11ax defines an *operating mode indication (OMI)* procedure for this purpose. As shown in Figure 2-12, the Wi-Fi 6 client that transmits a frame with an OM Control subfield is defined as the OMI initiator and the AP is the OMI responder. A Wi-Fi 6 client uses the OM Control subfield in 802.11 data and management frames to indicate a change of either transmission or receiver mode of operation.

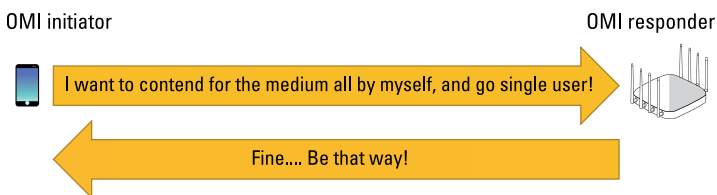


FIGURE 2-12: Transmit Operating Mode.

A client can switch between single-user or multi-user UL-OFDMA operations with a change in *transmit operating mode (TOM)*. Therefore, a Wi-Fi 6 client can both suspend and resume responses to the trigger frames sent by an AP during the UL-OFDMA process.

Additionally, a Wi-Fi client station can signal a change in *receive operating mode (ROM)* to the AP. The client indicates to the AP the maximum number of spatial streams and the maximum channel bandwidth that the client can support for downlink transmission. As shown in Figure 2-13, the client can indicate a change in channel size and number of supported spatial streams.

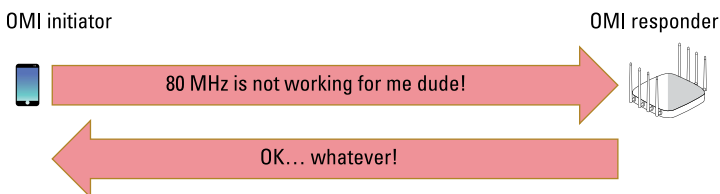


FIGURE 2-13: Receive Operating Mode.

You might surmise that OFDMA requires a lot of complex math, and you would be correct. Wi-Fi 6 APs require more processor power to perform the calculations needed for OFDMA operations. And legacy 802.11n/ac APs cannot be upgraded to perform OFDMA operations. Another question that I get all the time is “Are there any defined standards on how a Wi-Fi 6 AP makes the decisions on how to allocate the RUs to multiple clients?” The answer is no, and a lot of the RU allocation horsepower will depend on the radio chipset vendors. WLAN vendors may then further enhance their own airtime scheduling capabilities for RU allocation. As previously mentioned, RU allocation and scheduling mechanisms will likely become more precise in later generations of hardware. OFDMA is truly the secret sauce of Wi-Fi 6 that promises true multi-user communication.

- » Understanding MU-MIMO technology
- » Recognizing spatial diversity requirements
- » Applying MU-MIMO in the real world: PtMP bridging
- » Differentiating between MU-OFDMA and MU-MIMO

Chapter 3

HOLY COW, MU-MIMO

In this chapter, you learn about the other Wi-Fi 6 multi-user technology, MU-MIMO. All the first-generation Wi-Fi 6 radios support downlink MU-MIMO. Support for uplink MU-MIMO is being added to newer generations of Wi-Fi 6 hardware; however, support will be optional.

Introducing MU-MIMO

Wi-Fi 6 radios also support a secondary multi-user technology called *multi-user multiple-input multiple-output (MU-MIMO)*. The phrase is quite a mouthful. Some people like to pronounce MU-MIMO as *moo mahy-moh*, however, rest assured that the technology has nothing to do with cows. Much like OFDMA, MU-MIMO allows for multiple user communications downlink from an access point (AP) to multiple clients during the same transmission opportunity (TXOP). However, as opposed to partitioning the frequency space, MU-MIMO instead takes advantage of the fact that APs have multiple radios and antennas. A MU-MIMO AP transmits unique modulated data streams to multiple clients simultaneously (see Figure 3-1). The goal is to improve efficiency by using less airtime.




FIGURE 3-1: MU-MIMO – multi-user multiple-input multiple-output.



A five-number syntax is sometimes used when describing MU-MIMO radio capabilities. In a MU-MIMO system, the first number always references the transmitters (TX), and the second number references the receivers (RX). The third number represents how many unique single-user (SU) streams of data can be sent or received. The fourth number references how many multiple-user (MU) streams can be transmitted. A fifth number is used to represent a MU-MIMO group or how many MU-MIMO clients are receiving transmissions at the same time.

For example, when a MU-MIMO-capable AP operates using $4 \times 4:4:4:4$, four unique spatial streams would be destined to four independent MU-MIMO-capable clients (see Figure 3-1). However, when a MU-MIMO-capable AP operates as a $4 \times 4:4:4:2$ MU-MIMO AP, two unique spatial streams would be destined to one $2 \times 2:2$ client, and the other two spatial streams would be destined for a different $2 \times 2:2$ client (see Figure 3-2).




FIGURE 3-2: Downlink MU-MIMO – $4 \times 4:4:4:2$.

So how would this work if there are 20 MU-MIMO clients associated to the Wi-Fi 6 AP? The AP makes the decision as to which clients received the downlink MU-MIMO transmissions and which

clients are assigned to the MU-MIMO client group. For example, four clients could receive spatial streams simultaneously in the first downlink transmission and then four different clients might receive spatial streams simultaneously in the next downlink transmission.

Downlink MU-MIMO was first introduced in the second generation of 802.11ac radios. However, very few MU-MIMO-capable 802.11ac (Wi-Fi 5) clients were ever introduced to the marketplace, and the technology has rarely been used in the enterprise. Figure 3-3 displays the *maximum client capabilities* view within the ExtremeCloud IQ management platform. In this example, less than 10 percent of the clients support MU-MIMO. However, MU-MIMO capabilities of the client population will grow over time because Wi-Fi 6 radios are required to support downlink MU-MIMO.

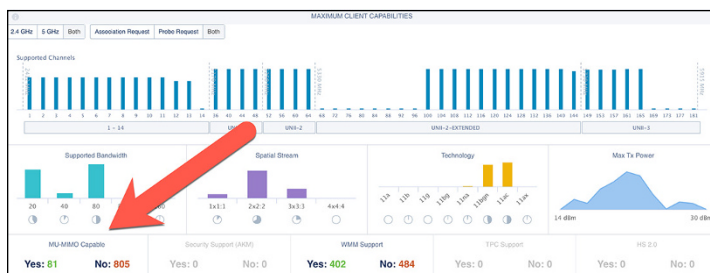


FIGURE 3-3: Maximum client capabilities.

A key difference between Wi-Fi 5 (802.11ac) MU-MIMO and Wi-Fi 6 MU-MIMO is how many MU-MIMO clients communicate with an AP at the same time. Wi-Fi 5 is limited to a MU-MIMO group of only four clients. Wi-Fi 6 is designed to support up to 8x8:8 MU-MIMO in both downlink and uplink, which allows it to serve up to eight users simultaneously and provide significantly higher data throughput.



MU-MIMO also requires *transmit beamforming (TxBF)*, which requires sounding frames. The sounding frames add excessive overhead, especially when the bulk of data frames are small. The overhead from the sounding frames usually negates any performance gained from a MU-MIMO AP transmitting downlink simultaneously to multiple 802.11ac clients. To address this issue, 802.11ax includes MU-MIMO enhancements like grouping

sounding frames, data frames, and other frames among multiple users to reduce overhead.



Downlink MU-MIMO is available in all Wi-Fi 6 radios. Support for uplink MU-MIMO was not available in the first generation of Wi-Fi 6 radios, however, vendors will have the option to support uplink MU-MIMO in second generation radios.

I Need Some Space

Wi-Fi 6 clients support downlink MU-MIMO, but MU-MIMO requires spatial diversity. Because of this, physical distance between the clients is necessary (See Figure 3-4). Additionally, MU-MIMO is more effective if the clients remain stationary. Moving targets are harder to hit! Even if all Wi-Fi clients support MU-MIMO, the majority of modern-day enterprise deployments of Wi-Fi involve a high density of users and devices, which are not ideal for MU-MIMO conditions.




FIGURE 3-4: Spatial diversity – MU-MIMO.

Almost all indoor WLANs are high-density (HD) environments because there are so many users and so many devices. Many of the users want to connect to an enterprise WLAN with as many as three or four Wi-Fi devices. Most high-density environments consist of multiple areas where roaming is also a top priority; therefore, clients are mobile and not stationary. The required spatial diversity

simply does not exist within the bulk of indoor enterprise Wi-Fi high-density deployments as depicted in Figure 3-5.




FIGURE 3-5: High-density enterprise Wi-Fi deployment.

Bridge over Troubled Waters

A very good use case for MU-MIMO is *point-to-multipoint (PtMP)* bridge links between buildings (see Figure 3-6). The spatial diversity that is required for MU-MIMO exists in this type of outdoor deployment. Bridge links require high bandwidth, which MU-MIMO can deliver with a PtMP deployment.




FIGURE 3-6: Point-to-multipoint (PtMP) bridge links.

What Is the Difference?

So how do the two Wi-Fi 6 multi-user technologies match up? Table 3-1 compares OFDMA and MU-MIMO.

TABLE 3-1 OFDMA and MU-MIMO Comparison

OFDMA	MU-MIMO
Increased efficiency	Increased capacity
Reduced latency	Higher data rates per user
Best for low-bandwidth applications	Best for high-bandwidth applications
Best with small packets	Best with large packets

MU-MIMO would theoretically be a favorable option in very low client density, high-bandwidth application environments where large packets are transmitted.



TECHNICAL
STUFF

The minimum RU size for MU-MIMO (downlink or uplink) is 106 subcarriers or greater.



REMEMBER

Wi-Fi allows for simultaneous use of both MU-OFDMA and MU-MIMO, but this is not expected to be widely implemented.

Do not confuse OFDMA with MU-MIMO. OFDMA enables multi-user access by subdividing a channel. MU-MIMO enables multi-user access by using different spatial streams.

IN THIS CHAPTER

- » Understanding medium contention overhead
- » Recognizing the cause of OBSS
- » Differentiating between basic service sets
- » Applying adaptive CCA thresholds

Chapter 4

A Splash of Wi-Fi Color

If you're reading a book, you might use highlighters to add color to specific important passages or to differentiate sections of text. Obviously, you can do the same to text in electronic documents and spreadsheets as well. Now, you can add color to your Wi-Fi! In this chapter, you learn about BSS color and spatial reuse operation, which have the potential to decrease medium contention overhead.

OBSS AKA Co-Channel Interference

Wi-Fi uses radio frequency communication, which is a *half-duplex medium* — only one radio can transmit on a frequency domain at any given time. A frequency domain is a fancy technical phrase for a channel. Everyone must take turns because if everyone “talks” at the same time, no data is communicated because no one is “listening.”



REMEMBER

Wi-Fi networks use the *carrier sense with multiple access collision avoidance (CSMA/CA)* method to ensure that only one radio can transmit on the same channel at any given time. An 802.11 radio defers transmissions if it hears the physical (PHY) preamble transmissions of any other 802.11 radio at a *signal detect (SD)* threshold of just four decibels or more above the noise floor.

CSMA/CA is necessary to avoid collisions; however, the deferral of transmissions also consumes valuable airtime. This problem is called *contention overhead*. Unnecessary medium contention overhead that occurs when too many access points (APs) and clients hear each other on the same channel is called an *overlapping basic service set (OBSS)*, shown in Figure 4-1. OBSS is also more commonly referred to as *co-channel interference*.

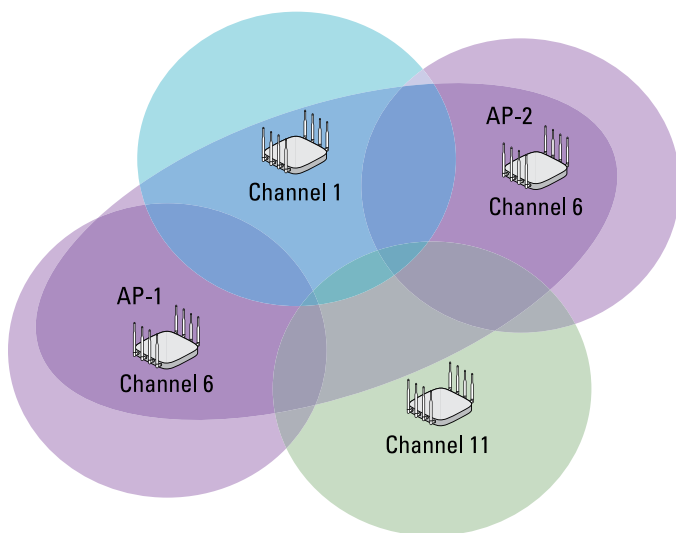


FIGURE 4-1: OBSS – Overlapping basic service set.

For example, if AP-1 on channel 6 hears the preamble transmission of a nearby AP (AP-2), also transmitting on channel 6, AP-1 defers and can't transmit at the same time. Likewise, all the clients associated to AP-1 must also defer transmission if they hear the preamble transmission of AP-2. The *basic service set (BSS)* is the cornerstone topology of Wi-Fi network. The communicating devices that make up a BSS consist of one AP radio with one or more client stations. OBSS creates medium contention overhead and consumes valuable airtime because you have two basic service sets on the same channel that can hear each other — thus, the term OBSS.

In reality, Wi-Fi clients are the primary cause of OBSS interference. As shown in Figure 4-2, if a client associated to AP-1 is transmitting on channel 36, it is possible that AP-2 (and any

clients associated to AP-2) will hear the PHY preamble of the client and must defer any transmissions.




FIGURE 4-2: OBSS interference caused by client.



REMEMBER

Due to the mobile nature of Wi-Fi client devices, OBSS interference isn't static. The situation changes as client devices move.

The primary goal of channel reuse patterns is to reduce co-channel interference (also known as OBSS). A channel reuse plan reduces airtime consumption caused by OBSS by isolating frequency domains. However, only three channels are used in the 2.4 GHz band. Because only three channels are available in the 2.4 GHz band and because OBSS is caused by clients, medium contention deferral is pretty much inevitable in the 2.4 GHz band. Co-channel interference can also be a problem in the 5 GHz band, especially if many of the 5 GHz channels are not available for a 5 GHz channel reuse plan. To increase capacity in dense environments, frequency reuse between basic service sets needs to be increased.

BSS Color

Wi-Fi 6 technology defines a method that may increase the channel reuse by a factor of eight. *BSS color*, also known as *BSS coloring*, is a method for addressing medium contention overhead due to OBSS. BSS color is an identifier of the basic service set (BSS).

In reality, the BSS color identifier is not a color, but instead is a numerical identifier. Wi-Fi 6 radios are able to differentiate between BSSs using a BSS color (numerical identifier) when other radios transmit on the same channel.



TECHNICAL
STUFF

BSS color information is communicated at both the PHY layer and the MAC sublayer. In the preamble of an 802.11ax PHY header, the SIG-A field contains a 6-bit BSS color field. This field can identify as many as 63 BSSs. At the MAC sublayer, BSS color information is seen in 802.11 management frames. The HE operation information element contains a subfield for BSS color information. Six bits can be used to identify as many as 63 different colors (numerical values) and represent 63 different BSSs. The goal is for Wi-Fi 6 radios to differentiate between BSSs using a BSS color identifier when other radios transmit on the same channel.



REMEMBER

BSS color detects a color bit in the PHY header of a Wi-Fi 6 radio frame transmission. This means that legacy 802.11a/b/g/n radios will not be able to interpret the color bits because they use a different PHY header format.

As shown in Figure 4-3, the purpose of BSS color is to uniquely identify different BSSs even though they are transmitting on the same channel. Keep in mind that this figure is a visual illustration and that the color information is actually a numerical value. When a Wi-Fi 6 radio is listening to the medium and hears the PHY header of an 802.11ax frame sent by another Wi-Fi 6 radio, the listening radio will check the BSS color of the transmitting radio. Channel access is dependent on whether the BSS color is the same or different as described below:

- » **Intra-BSS:** If the color is the same, then the frame is considered an *intra-BSS* transmission, and the listening radio will defer. If the color is the same, this is considered to be an intra-BSS frame transmission. In other words, the transmitting radio belongs to the same BSS as the receiver; therefore, the listening radio will defer.
- » **InterBSS:** If the color is different, then the frame is considered an *inter-BSS* transmission from an OBSS. In other words, the transmitting radio belongs to a different BSS and deferral may not be necessary for the listening radio.




FIGURE 4-3: BSS color.

Spatial Reuse Operation

Using a procedure called *spatial reuse operation (SRO)*, Wi-Fi 6 can apply adaptive clear channel assessment (CCA) thresholds for detected OBSS frame transmissions. The goal of BSS color and spatial reuse is to ignore transmissions from an OBSS and therefore be able to transmit at the same time. Every device can calm down and focus on its own job.

In the example shown in Figure 4-4, three BSSs transmitting on channel 36 have been assigned the BSS colors of red, blue, and green. This would be considered an inter-BSS environment. The APs and clients that belong to the red and green BSSs are within close proximity of each other and deferral is most likely needed. However, despite being on the same channel, deferral may not be necessary between the APs and clients that are members of the red and blue BSSs. Because of the greater physical distance, an adaptive CCA threshold might be used.

802.11 (Wi-Fi) radios use a *clear channel assessment (CCA)* to appraise the RF medium. If the RF medium is busy, an 802.11 radio does not transmit and instead defers for a period of time, called a *slot time*. The CCA involves listening for RF transmissions at the Physical layer. 802.11 radios use two separate CCA thresholds when listening to the RF medium. The *signal detect (SD)* threshold is used to identify the incoming 802.11 preamble transmission from another transmitting 802.11 radio. The preamble is a

component of the Physical layer header of 802.11 frame transmissions. The signal detect (SD) threshold is statistically close to a 4 dB signal-to-noise ratio (SNR) for most 802.11 radios to detect and decode an 802.11 preamble. In other words, an 802.11 radio can usually decode any incoming 802.11 preamble transmissions at a received signal at about 4 dB above the noise floor.




FIGURE 4-4: Inter-BSS.

The *energy detect (ED)* threshold is used to detect any other type of RF transmissions during the clear channel assessment (CCA). The ED threshold is 20 dB higher than the signal detect threshold. Think of the signal detect threshold as a method of detecting and deferring Wi-Fi radio transmissions. Think of the energy detect threshold as a method of detecting and deferring any signals from non-802.11 transmitters. Both thresholds are used together during the CCA to determine whether the medium is busy and therefore must defer transmissions.

OBSS interference is a result of radios deferring based on the signal detect (SD) threshold being so low. Statistically, most radios can decode an 802.11 preamble if the received signal is only 4 dB above the noise floor. As a result of this very low signal detect threshold, APs and clients on the same channel hear each other and will defer, despite being separated by significant physical distance.

Many enterprise WLAN vendors offer a configuration threshold setting on APs called *receive start of packet (RX-SOP)*. This configuration setting gives administrators the capability to fine-tune and manipulate the signal detect threshold of APs. For example, an admin could manually set the SD threshold to a less sensitive

value as opposed to 4 dB above the noise floor. A less sensitive SD threshold will result in a reduction in co-channel interference from nearby APs and clients on the same channel. The RX-SOP setting should not be taken lightly. If the threshold is raised too high, transmissions from nearby APs on the same channel will corrupt data and cause a degradation of performance that is actually worse than the performance loss due to the deferral of OBSS interference.

Spatial reuse operation (SRO) allows Wi-Fi 6 radios to apply adaptive clear channel assessment (CCA) signal detect thresholds. Think of SRO as a dynamic implementation of the static RX-SOP threshold settings.

Based on the detected BSS color, Wi-Fi 6 radios can implement an adaptive CCA that can raise the signal detect threshold for inter-BSS frames, while maintaining a lower threshold for intra-BSS traffic. If the signal detect threshold is raised higher for incoming OBSS frames, a radio might not need to defer, despite being on the same channel. The adaptive signal detect threshold can be adjusted on a per-color and per-frame basis for inter-BSS traffic. In the example in Figure 4-5, a static SD threshold of -96 dBm (decibels relative to 1 milliwatt) might be used for the reception of intra-BSS traffic, whereas an adaptive SD threshold between -96 dBm to -83 dBm might be used for inter-BSS traffic.




FIGURE 4-5: Spatial reuse operation - Adaptive CCA.



BSS color together with spatial reuse operation has the potential to decrease the OBSS channel contention problem that is symptomatic of existing low SD thresholds. Will BSS color solve the co-channel interference problem? The answer is: Probably not any time soon. First, keep in mind that legacy clients do not have the capability to distinguish between BSSs on the same channel. Second, getting the adaptive SD thresholds to work properly will

be a challenge for the Wi-Fi 6 radio chipset manufacturers. Much as RX-SOP can be a dangerous configuration setting, the dynamic and adaptive SD thresholds could also have a negative impact. While BSS color and SRO hold potential, the technology may take a long time to mature before it is effective in a real-world Wi-Fi environment.

- » Conserving power with TWT
- » Achieving higher data speeds with 1024-QAM
- » Defining new PHY headers
- » Utilizing 20 MHz-only mode

Chapter 5

Additional Wi-Fi 6 Enhancements

You can never have too many good things, can you? Wi-Fi 6 is overflowing with brand new features that make life better. In this chapter, you learn about several other Wi-Fi 6 enhancements including target wake time (TWT), 1024 quadrature amplitude modulation (1024-QAM), and new frame formats, as well as other Wi-Fi 6 efficiency mechanisms.

Target Wake Time (TWT)

Target wake time (TWT) is a Wi-Fi 6 enhanced power-saving mechanism. A TWT is a negotiated agreement, based on expected traffic activity between the access point (AP) and clients, to specify a scheduled target wake-up time for Wi-Fi 6 clients in power-save (PS) mode. The negotiated TWTs allow an AP to manage client activity by scheduling client stations to operate at different times in order to minimize contention between the clients. A TWT reduces the required amount of time that a client station in PS mode needs to be awake.



This enhancement allows the client to “sleep” longer and reduce energy consumption. As opposed to legacy client power-saving mechanisms such as delivery traffic indication map (DTIM), which require sleeping client devices to wake up in microsecond intervals, TWT could theoretically allow client devices to sleep for hours. TWT is thus an ideal power-saving method for mobile devices and Internet of Things (IoT) devices that need to conserve battery life.

As depicted in Figure 5-1, a TWT frame exchange is used between the AP and the clients to negotiate a scheduled TWT. For each Wi-Fi 6 client, there can be as many as eight separate negotiated scheduled wake-up agreements for different types of application traffic. Once the negotiation is complete, the clients sleep and then awaken at the targeted intervals. Wi-Fi 6 has also extended TWT functionality to include a non-negotiated TWT capability. An AP can create wake-up schedules and deliver TWT values to the Wi-Fi 6 clients via a broadcast TWT procedure.

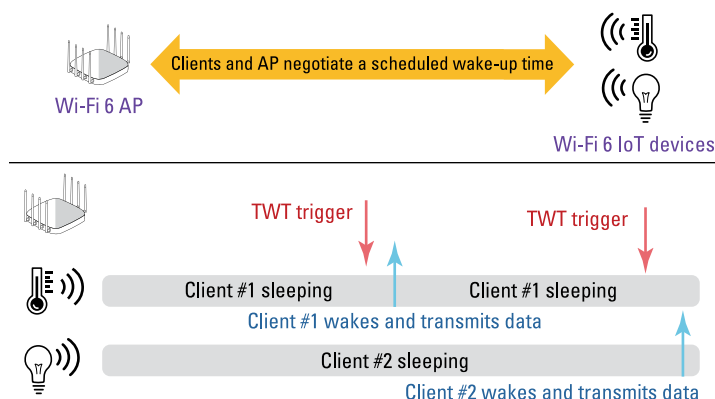


FIGURE 5-1: Target wake time.

The TWT power-saving capabilities are ideal for Internet of Things (IoT) devices. TWT was originally defined in the 802.11ah amendment, which defined the use of Wi-Fi in frequencies below 1 GHz. Despite the goal of moving IoT devices to a lower frequency band, most IoT devices with a Wi-Fi radio still currently transmit in the 2.4 GHz frequency band. Because 802.11ax now also defines TWT, these same extended power-saving capabilities could be available to IoT devices with 802.11ax radios that transmit in the 2.4 GHz band. It remains to be seen if IoT device companies will manufacture IoT devices with 802.11ax radios; however, the potential to conserve battery life is attractive if the devices are TWT-capable.

1024-QAM

Although the primary goal of Wi-Fi 6 is increased efficiency, more speed is not a bad thing. Elevated efficiency and more speed are not mutually exclusive goals. *Quadrature amplitude modulation* (QAM) uses both the phase and amplitude of a radio frequency signal to represent data bits. Wi-Fi will support 1024-QAM and new modulation and coding schemes (MCSs) that define higher data rates.



TECHNICAL
STUFF

For comparison, 256-QAM (introduced in 802.11ac) modulates 8 bits per symbol, whereas 1024-QAM modulates 10 bits per symbol — a potential 25 percent increase in data throughput. Wi-Fi 6 also introduces two new MCSs that make use of 1024-QAM modulation: MCS-10 and MCS-11, both of which are optional. In first generation Wi-Fi 6 radios, 1024-QAM may only be used with 242-subcarrier resource units (RUs) or larger. This means that at least a full 20 MHz of channel bandwidth will be needed for 1024-QAM.

Much like 256-QAM, very high signal-to-noise ratio (SNR) thresholds (35 decibels or more) will be needed in order for Wi-Fi 6 radios to use 1024-QAM modulation. Pristine radio frequency environments with a low noise floor and close proximity between a Wi-Fi 6 client and a Wi-Fi 6 AP will most likely be needed.

A constellation diagram, also known as a *constellation map*, is a two-dimensional diagram often used to represent QAM modulation. A constellation diagram is divided into four quadrants, and different locations in each quadrant can be used to represent data bits. Areas on the quadrant relative to the horizontal axis can be used to represent various phase shifts. Areas relative to the vertical axis are used to represent amplitude shifts.

The number of points in the modulation constellation map determines the number of bits conveyed with each symbol. Figure 5-2 shows a comparison of constellation charts between 256-QAM and 1024-QAM modulation. As you can see, 1024-QAM has many more constellation points. *Error vector magnitude (EVM)* is a measure used to quantify the performance of a radio receiver or transmitter with regard to modulation accuracy. With QAM modulation, EVM is a measure of how far a received signal is from a constellation point. Any Wi-Fi 6 radios that use 1024-QAM modulation will need strong EVM and receive sensitivity capabilities.




FIGURE 5-2: 256- and 1024-QAM.

Long Symbol Time and Guard Intervals

For digital signals, data is modulated onto the carrier signal in bits or collections of bits called *symbols*. 802.11ax introduces a longer OFDM *symbol time* of 12.8 μs (microseconds), which is four times the legacy symbol time of 3.2 μs . The increase in the number of subcarriers (tones) also increases the OFDM symbol duration. Subcarrier spacing is equal to the reciprocal of the symbol time. The symbol time used is four times longer, as 802.11ax uses subcarrier spacing of 78.125 KHz, which is one quarter the size of legacy 802.11n/ac subcarrier spacing.

The *guard interval* (GI) is a period of time between symbols that accommodates the late arrival of symbols over long paths. In a multipath environment, symbols travel different paths, so some symbols arrive later. A “new” symbol may arrive at a receiver before a “late” symbol has been completely received. This problem is known as *intersymbol interference* (ISI) and can result in data corruption. The *delay spread* is the time differential between multiple paths of the same signal. Normal delay spread is from 50 nanoseconds to 100 nanoseconds, and a maximum delay spread is about 200 nanoseconds. The guard interval should be two to four times the length of the delay spread. Think of the guard interval as a buffer for the delay spread.

802.11a/g defined the use of a 0.8 μs (which equals 800 nanoseconds) guard interval, while 802.11n/ac also added the option for a 0.4 μs (400 nanosecond) short guard interval, which was intended

for use in indoor environments. When the legacy symbol time of $3.2\ \mu\text{s}$, which is used for the modulated data, is combined with the standard $0.8\ \mu\text{s}$ guard interval, the total symbol duration is $4.0\ \mu\text{s}$. When the legacy data symbol time of $3.2\ \mu\text{s}$ is combined with the $0.4\ \mu\text{s}$ short guard interval, the total symbol duration is $3.6\ \mu\text{s}$.

Wi-Fi 6 radios utilize three different guard intervals that can be used together with the $12.8\ \mu\text{s}$ symbol time that is used for the modulated data:

- » **0.8 μs Guard Interval:** This guard interval is likely to be used for most indoor environments. When combined with the time of $12.8\ \mu\text{s}$ that is used for the data, the total symbol time for indoor communications will be $13.6\ \mu\text{s}$.
- » **1.6 μs Guard Interval:** This guard interval is intended for outdoor communications. When combined with the time of $12.8\ \mu\text{s}$ that is used for the data, the total symbol time will be $14.4\ \mu\text{s}$. The guard interval may be needed in high multipath indoor environments to ensure the stability of uplink OFDMA or uplink MU-MIMO communication.
- » **3.2 μs Guard Interval:** This guard interval is also intended for outdoor communications. When combined with the time of $12.8\ \mu\text{s}$ that is used for the data, the total symbol time will be $16.0\ \mu\text{s}$. The longer symbol time and longer guard intervals will provide for more robust outdoor communications.

New PHY Headers

Wi-Fi 6 adds a new physical (PHY) header to all 802.11 frames. The PHY header contains a preamble and other information used for the initial communications setup between two radios. As shown in Figure 5-3, Wi-Fi 6 radios make use of four new PHY headers to support high efficiency (HE) radio transmission:

- » **HE SU:** The high efficiency single-user PHY header is used for single-user transmissions.
- » **HE MU:** The high efficiency multi-user PHY header is used for transmissions to one or more users. This format is not used

as a response to a trigger, which means this PHY header is used for trigger frames or downlink transmissions.

- » **HE ER SU:** The high efficiency extended-range single-user format is intended for a single user. Portions of this PHY header are boosted by 3 decibels to enhance outdoor communications and range.
- » **HE TB:** The high efficiency trigger-based format is for a transmission that is a response to a trigger frame. In other words, this PHY header format is used for uplink communications.




FIGURE 5-3: Wi-Fi 6 – PHY header formats.



The preamble is used for synchronization between transmitting and receiving radios and consists of two parts: legacy and high efficiency (HE). The legacy preamble is easily decodable by legacy stations (STAs) and is included for backward compatibility. The HE preamble components are used to communicate information between Wi-Fi 6 radios about OFDMA, MU-MIMO, BSS color, and more.

20 MHz-Only Mode

Some Wi-Fi 6 client radios can take advantage of a *20 MHz-only* mode of operation. Client stations will be able to inform an AP that they are operating as 20 MHz-only clients. As shown in Figure 5-4, a 20 MHz-only client can still operate within a 40 MHz, 80 MHz, or 160 MHz channel. However, the 20 MHz-only clients must communicate via RUs of the *primary channel*.




FIGURE 5-4: 20 MHz-only Wi-Fi 6 Client.

Very specific 26-tone RU, 52-tone RU, or 106-tone RU mappings within the primary channel are used by a 20 MHz-only client. This ensures that a 20 MHz-only client is only assigned the proper OFDMA tone mappings and RU allocations that the client can support even if larger channels are being used. As opposed to a smartphone or laptop client device, 20 MHz-only clients are small form factor devices with limited processing capability and lower power requirements. The 20 MHz-only operational mode is ideal for IoT clients that could take advantage of the Wi-Fi 6 power-saving capabilities but not necessarily need the full capabilities that Wi-Fi 6 has to offer. This configuration allows client manufacturers to design less complex chipsets at a lower cost, which is ideal for IoT devices.

Multi-TID AMPDU

Two terms that everyone should understand are MSDU and MPDU. An 802.11 MAC *Service Data Unit (MSDU)* is the layer 3–7 payload of an 802.11 data frame. An 802.11 MAC *Protocol Data Unit (MPDU)* is essentially a technical term for a wireless frame. An MPDU consists of a frame header, body, and trailer with the MSDU payload encapsulated in the frame body.

Frame aggregation is a method of combining multiple frames into a single frame transmission. Fixed MAC layer overhead and medium contention overhead are reduced, which results in less airtime consumption. The most common method of frame aggregation is known as *aggregate MAC protocol data unit (A-MPDU)*. Multiple MPDUs can be aggregated into a single transmission. A-MPDU comprises multiple MPDUs and is prepended with a PHY header.

Prior to Wi-Fi 6, all the individual MPDUs must have been of the same 802.11e QoS access category when A-MPDU frame aggregation was used. Voice MPDUs could not be mixed with Best Effort or Video MPDUs within the same aggregated frame.

As shown in Figure 5-5, Wi-Fi 6 introduces *multi-traffic identifier aggregated MAC protocol data unit (multi-TID AMPDU)*, which allows the aggregation of frames from multiple traffic identifiers (TIDs), from the same or different QoS access categories. The capability to mix MPDUs of different QoS traffic classes allows Wi-Fi 6 radios to aggregate more efficiently, reducing overhead and thus increasing throughput and therefore overall network efficiency.




FIGURE 5-5: Multi-TID AMPDU.

- » Celebrating the spectrum bonanza and worldwide adoption
- » Defining 6 GHz device classes and regulations
- » Exploring Automated Frequency Coordination
- » Understanding discovery mechanisms
- » Deploying channel reuse best practices

Chapter 6

Wi-Fi 6E – A New Beginning in 6 GHz

Wi-Fi is setting up shop in the 6 GHz frequency band. Using Wi-Fi in this pristine RF environment is a wireless networking game-changer. In this chapter, you will learn about Wi-Fi 6E operations and real-world operational considerations.

Spectrum Bonanza

In early 2020, the U.S. Federal Communications Commission (FCC) voted unanimously to make 1,200 megahertz of spectrum in the 6 GHz band available for unlicensed use in the United States. To put this in perspective, the new 6 GHz spectrum available for Wi-Fi is more than double the usable channels of the 2.4 GHz and 5 GHz channels combined. So effectively, it triples the available unlicensed spectrum available for Wi-Fi. This, my friends, is an enormous spectrum bonanza.

As shown in Figure 6-1, in the United States, there are as many as 59 new 20 MHz channels available across four U-NII bands. Additionally, this new availability includes the potential to use 29 new 40 MHz channels. And yes, I am expecting the 14 new 80 MHz band to be used extensively in the enterprise. All of these channels are defined by a center frequency and have a channel number for identification.




FIGURE 6-1: 6 GHz U-NII bands and Wi-Fi channels (United States).

In late 2020, the Wi-Fi Alliance announced Wi-Fi 6E as an extension for certifying the 802.11ax features and capabilities of Wi-Fi 6 to the 6 GHz band. Wi-Fi 6E is the industry marketing name that identifies Wi-Fi 6 devices that operate in 6 GHz.

One key difference of using the 6 GHz frequency band for 802.11ax technology is there is no need for backward compatibility. Because 802.11a/b/g/n/ac radios operate only on the 2.4 GHz or 5 GHz band, and not the 6 GHz band, protection mechanisms aren't needed. The 6 GHz frequency band will be a "pure" 802.11ax technology band for Wi-Fi communication (at least, as pure as anything on the internet can be). The 6 GHz band has no dynamic frequency selection (DFS) requirements needed to avoid radar. Some incumbent transmitters in the 6 GHz band require a different type of interference protection, which will be discussed later in this chapter.

WLAN vendors are already selling Wi-Fi 6E access points (APs) in many different form factors. In most cases, these devices have radios for all three bands (2.4, 5, and 6 GHz). However, only new Wi-Fi 6E client devices with 6 GHz radios can communicate

with the 6 GHz radio in a Wi-Fi 6E AP. An older dual-frequency (2.4 and 5 GHz) smartphone only communicates with the 2.4 or 5 GHz radio in a tri-band Wi-Fi 6E AP.

6 GHz Worldwide

As shown in Figure 6-2, many world regions are making all or portions of the 6 GHz frequency band available for Wi-Fi. As of this writing, 50 countries have approved or are considering new regulations for the unlicensed use of 6 GHz. The Wi-Fi Alliance maintains a web page with a current list of countries enabling Wi-Fi in the 6 GHz band at www.wi-fi.org/countries-enabling-wi-fi-6e. Ratification for unlicensed use usually follows regulatory approval in a timely fashion.

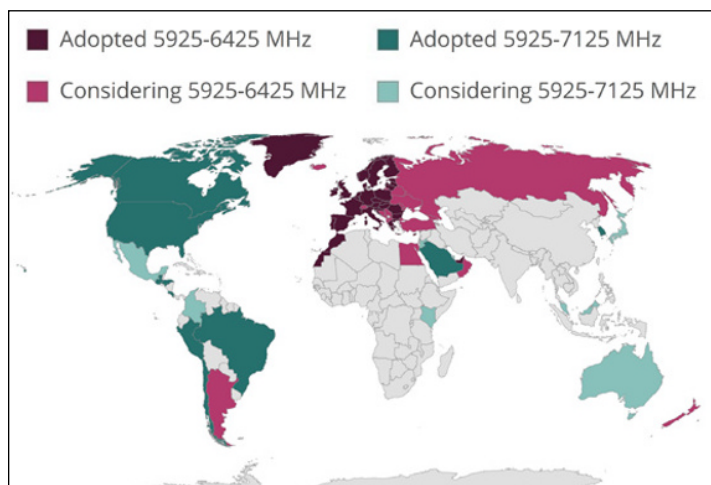


FIGURE 6-2: The worldwide view of 6 GHz Wi-Fi.

For example, Canada and the European Union (EU) already approved the reservation of the spectrum. Canada expects to ratify in late fall to early 2022, and EU ratification of Wi-Fi 6E operation is expected sometime after that. Some regions will be more restrictive than others. For example, only 500 MHz of 6 GHz frequency space will be available for Wi-Fi 6E in the EU. As shown in Figure 6-3, European regulators are focusing only on the UNII-5 band (5.925–6.425 GHz). This means that 24 new 20 MHz channels, 12 new 40 MHz

channels, and six new 80 MHz channels are available for Wi-Fi 6E communications in Europe.




FIGURE 6-3: 6 GHz Wi-Fi in Europe.

Device Classes and Regulations

Each country defines power regulations on how these 6 GHz bands can be used for Wi-Fi. For this book, I will focus on the United States regulations as mandated by the FCC. The good news is that the international regulatory organizations for spectrum management communicate with one another. Therefore, many of the rules and regulations outlined in this book will be very similar in other regions of the world.



REMEMBER

The FCC is authorizing the use of the entire 6 GHz band and all four of the U-NII bands for indoor use. However, the FCC does define seven different classes of Wi-Fi equipment that can be used. The rules of operation are different for these seven different types of devices. To keep it simple, I focus on the differences between indoor and outdoor Wi-Fi regulations.

As shown in Table 6-1, a new device class of *low-power indoor (LPI)* APs is allowed to transmit indoors only with a maximum

equivalent isotropically radiated power (EIRP) of 30 dBm. Additionally, low-power indoor clients can connect to the LPI APs with a maximum EIRP of 24 dBm. Client devices in Wi-Fi 6E, for both low power and standard power operations, are restricted to 6 dB lower than the AP.

TABLE 6-1 Expanded unlicensed use of the 6 GHz band (United States)

Device Class	Operating Band	Maximum EIRP
Low-Power Indoor (LPI) AP	U-NII-5 (5.925-6.425 GHz)	30 dBm
Client Connected to LPI AP	U-NII-6 (6.425-6.525 GHz)	24 dBm
	U-NII-7 (6.525-6.875 GHz)	
	U-NII-8 (6.875-7.125 GHz)	
Standard-Power AP (AFC Controlled)	U-NII-5 (5.925-6.425 GHz)	36 dBm
Client Connected to Standard-Power AP	U-NII-7 (6.525-6.875 GHz)	30 dBm

The FCC will also require that all low-power devices incorporate permanently attached integrated antennas. Requiring an integrated antenna makes it significantly more difficult for someone to replace a device’s antenna with a higher gain antenna. The FCC has determined that abiding by these indoor power restrictions will not interfere with any incumbent outdoor services. As a result, effectively, the entire 6 GHz band will be available for indoor Wi-Fi in the United States. Some of the other regulations for LPI devices include:

- » Low-power indoor APs are limited to indoor locations, have an integrated antenna, and cannot use a weatherized enclosure.
- » Low-power indoor APs are prohibited on oil platforms, cars, trains, boats, and aircraft. The exception is that LPI APs can be used for in-flight entertainment systems in large passenger aircraft above 10,000 feet.
- » Low-power indoor APs are prohibited for control of or communications with unmanned aircraft systems. So sorry, no 6 GHz Wi-Fi drones.

- » Low-power indoor APs must be powered by a wired connection and not by battery power. LPI APs may use battery backup only during power outages.
- » Client devices must remain indoors while under the control of the LPI AP. Indoor clients are prohibited from making a direct air interface connection to other clients.

In the United States, the rules for outdoor Wi-Fi in the 6 GHz band are quite different from what is permitted indoors. As indicated earlier in Table 6-1, the U-NII-6 and U-NII-8 bands are *not* available at all for unlicensed outdoor communications. The U-NII-6 and U-NII-8 bands are already licensed for mobile satellite services and *fixed satellite services* (FSS) used in the broadcast and cable industries. Therefore, these two bands will be unavailable for outdoor Wi-Fi. The FCC does define another device class of a *standard-power AP* for unlicensed outdoor communications in the U-NII-5 and U-NII-7 bands. The maximum EIRP for a standard-power AP is 36 dBm. Clients that connect to a standard-power AP can have a maximum EIRP of 30 dBm. Additionally, expect spectrum management restrictions to protect licensed incumbent fixed services in the U-NII-5 and U-NII-7 bands.

Automated Frequency Coordination

For outdoor Wi-Fi communications in the U-NII-5 and U-NII-7 bands, the FCC will mandate the use of *automated frequency coordination* (AFC) to protect the incumbents. An AFC system will use geolocation databases to manage real-time frequency assignments to protect incumbent operations from RF interference. Before transmitting, a standard-power AP is required to obtain a list of permissible frequencies or a list of prohibited frequencies on which it cannot transmit, from an AFC system. The geographic coordinates of the AP should be automatically determined by GPS or a similarly reliable method before checking in with the AFC system.

For example, you might want to deploy an outdoor standard-power AP in Brookhaven, Georgia. Using either GPS or another method, you determine that the latitude of the AP is 33.865105 N, longitude is 84.336594 W, and the elevation is 304 meters. The AP would then use automated mechanisms to register its coordinates

with an FCC-approved AFC system provider. The AP's antenna height is also taken into consideration. The database of the AFC system provider automatically checks for possible interference with any incumbents. The AFC system provider determines the exclusion zones where standard power APs might cause harmful interference to incumbent links in the U-NII-5 and U-NII-7 bands. If an incumbent *fixed service (FS)* is nearby, a three-dimensional AFC protection contour is enforced on the standard-power Wi-Fi AP.

As shown in Figure 6-4, an AFC system provider calculates a protection area around every incumbent fixed service (FS) receiver using licensee data in the FCC's Universal Licensing System (ULS). Standard-power APs operating above 5 dBm/MHz and all outdoor APs must send their 3D location information to the AFC before they are allowed to transmit. The AFC system provider uses incumbent protection contours that account for an AP's geolocation, power, and other variables to determine which channels are permissible for the AP to transmit.




FIGURE 6-4: Automated Frequency Coordination.

In simpler words, the AP may not be allowed to transmit or may be required to lower its power well below 36 dBm EIRP to avoid interference with the incumbent. If there are no nearby incumbents in the area, the AP can transmit. AFC will be essential in urban areas and large cities where there are numerous incumbents. Think of this system like air traffic control to keep planes from crashing into each other. The automatic mechanisms just don't need time off, coffee, or a fancy tower near the airport.

Although the FCC has used coordinated spectrum management systems in other frequency bands with success, all of the AFC rules and system providers are not yet in place. In most countries, low-power 6 GHz APs for indoor-only use will hit the marketplace before standard-power outdoor APs are available in the commercial market. Because so many incumbents exist in the United States, final approval for AFC implementation is not expected until late 2022 or even 2023. We may see the outdoor deployment of standard power APs in other countries sooner.



TECHNICAL
STUFF

There has also been some discussion about allowing the use of *very low power (VLP)* devices to operate across the entirety of the 6 GHz band (5.950–7.125 GHz), both indoors and outdoors, without using an AFC. These devices are intended for high-speed (transportation) or short-range client-to-client devices communication. Without getting too detailed, the power limits would translate to a maximum EIRP of 5 dBm for a 20 MHz channel VLP device. Currently, the FCC is prohibiting the operation of VLP devices in the United States. However, approval of VLP devices is expected in other countries.

The Road to 6 GHz AP Discovery

Wi-Fi clients have traditionally used an active hunt-and-seek method to scan for APs. Clients send out probe request frames across the 2.4 and 5 GHz channels to discover APs. In 6 GHz, this traditional active scanning method is no longer efficient for initial AP discovery and even worse for roaming between APs. Client probing simply takes too much time because there are so many channels in the 6 GHz band. Wi-Fi clients can only send probe requests on 20 MHz channels, and scanning all 59 of the 20 MHz channels is not an option. Passive scanning on each channel would take even longer. It would take over 6 seconds for a client to listen for AP beacon frames on all 59 channels. Six seconds is an eternity in Wi-Fi communications. Therefore, new in-band and out-of-band AP discovery mechanisms have been designated for Wi-Fi 6E clients.

Surprisingly, the out-of-band discovery processes are expected to be the most widely used, even for Wi-Fi 6E clients already associated with a 6 GHz AP radio. Most chipsets used in Wi-Fi 6E client radios will also have 2.4 and 5 GHz capabilities, meaning they can

also scan and connect to APs using the legacy frequency bands. A tri-band AP can inform a Wi-Fi 6E client actively probing the 2.4 GHz or 5 GHz bands about the existing 6 GHz radio co-located in the AP. Therefore, there are two defined out-of-band discovery methods:

- » Reduced neighbor report (RNR)
- » Multiple BSSID beacon frames

802.11v first defined the possible use of a *reduced neighbor report* (RNR) information element that can be used to include information about a neighbor AP. For Wi-Fi 6E, the “neighbor AP” is actually the 6 GHz radio housed in the same AP along with the 2.4 GHz and 5 GHz radios. Wi-Fi 6E clients will learn about the available 6 GHz radio from the RNR information in either beacon or probe response frames sent by the AP’s 2.4 and 5 GHz radios.

In the example shown in Figure 6-5, a Wi-Fi 6E client sends directed probe requests across the 5 GHz band for an SSID called blue. Three APs answer with probe responses that carry basic service set (BSS) parameters for the blue SSID for the 5 GHz channels of 36, 40, and 44. However, inside each probe response is also RNR information about the 6 GHz radios transmitting on channels 53, 85, and 117. The SSID might be the same across bands or different in each frequency. The client can then decide whether to connect to 5 GHz AP radio, or even more likely, the available 6 GHz AP radio. Obviously, the goal is to eliminate probing time on the 6 GHz band. The client device can be informed of available 6 GHz BSSs without ever scanning the 6 GHz band. These devices get a better roadmap, or maybe just an easier-to-read menu.




FIGURE 6-5: Out-of-band discovery.



The same out-of-band discovery method is used when the Wi-Fi 6E client probes the 2.4 GHz band. The probe responses from the 2.4 GHz radios in the APs will respond about 2.4 GHz channel availability and RNR information about the 6 GHz radios co-located in the same AP. Figure 6-6 shows an RNR information element from a 5 GHz probe response frame. Note that it indicates a 6 GHz primary channel of 37 that is co-located in the same AP. The *operating class* is an indication of the 6 GHz channel size. An operating class of 134 indicates a 160 MHz channel. Likewise, 133 denotes an 80 MHz channel, 132 indicates a 40 MHz channel, and 131 indicates a channel size of 20 MHz. The reduced neighbor report (RNR) can also show whether or not the 6 GHz SSID is the same as the 5 GHz SSID. And the *Short SSID* parameter is effectively a hash of the 6 GHz SSID.

Information Element	Details
✓ Reduced Neighbor Report	Operating Class: 134, Channel: 37, BSSID 0: 6C:CD:D6:1D:01:0D, Short SSID 0: 5F473EF7
Element ID:	201
Length:	17 bytes
> TBTT Information Header:	0x0d00
Operating Class:	134
Channel Number:	37 ← 6 GHz channel info in a 5 GHz probe response
TBTT 0	
Neighbor AP TBTT Offset:	Unknown 0xff
BSSID:	6C:CD:D6:1D:01:0D
Short SSID:	5F473EF7
BSS Parameters:	0x6e
... ..0	OCT Recommended: No
... ..0.	Same SSID: No
... ..1..	Multiple BSSID: Yes
... ..1...	Transmitted BSSID: Yes
... ..0	Member of ESS with 2.4/5 GHz Co-Located AP: No
... ..1.	Unsolicited Probe Responses Active: Yes
... ..1.	Co-Located AP: Yes

FIGURE 6-6: Reduced neighbor report.

So, what if a client connects to a 6 GHz AP on channel 53 and wants to roam to another 6 GHz AP? Believe it or not, the most likely client active scanning method will once again be for the Wi-Fi 6E client to probe the 2.4 and/or 5 GHz channels to get RNR information about possible nearby 6 GHz APs to which the client might roam.

Another defined out-of-band discovery method utilizes *multiple BSSID* beacon frames and probe responses. Multiple BSSID is also a capability that was originally specified in the IEEE 802.11v amendment. It reduces management frame overhead by eliminating the need for multiple beacons for multiple SSIDs and BSSIDs. For example, the SSID/BSSID information for the three SSIDs of employees, guests, and voice could be consolidated into a single

beacon frame. Although this 802.11v capability has not been leveraged in the past, Wi-Fi 6E clients may take advantage of it to passively learn about multiple SSIDs/BSSIDs available across the multiple bands housed in one AP.



TIP

Out-of-band discovery should be the primary method that Wi-Fi 6E clients use to discover 6 GHz APs. However, there are also three potential methods for in-band discovery of 6 GHz APs:

- » **Passive:** Fast Initial Link Setup (FILS) discovery announcement frames
- » **Passive:** Unsolicited probe response frames
- » **Active:** Preferred Scanning Channels (PSC)

The first two in-band (6 GHz) discovery methods are passive. A *FILS discovery announcement frame* is analogous to a condensed beacon. Only critical information such as SSID, BSSID, and channel can be found in a FILS frame. If implemented, the 6 GHz AP sends this broadcast action frame out every 20 time units (TUs), approximately 20 milliseconds. The second passive method uses *unsolicited probe response frames*. A probe response frame contains all the same detailed information as a beacon. If implemented, the 6 GHz AP sends these unsolicited probe responses to a broadcast address approximately every 20 milliseconds. You will note that both passive methods are timed well for VoWiFi mobile devices that might use off-channel passive scanning. Both methods are also referred to as *fast passive scanning*.

The third in-band AP discovery method is an active method and the only way in which Wi-Fi 6E clients can send probe requests. Active scans on a *preferred scanning channel (PSC)* are the only method used for in-band probing. Effectively, Wi-Fi 6E clients can only send probe requests on every fourth 20 MHz channel. As shown in Figure 6-7, PSC channels also serve as the primary channels when channel bonding is used for 80 MHz channels. The complete list of PSC channels is 5, 21, 37, 53, 69, 85, 101, 117, 133, 149, 165, 181, 197, 213, and 229. Once again, the goal is that clients do not have to probe across all 59 channels. In fact, all three of these in-band discovery methods were originally intended for APs that *only* have a 6 GHz radio. In the real world, most vendors will manufacture tri-frequency band APs. I expect PSC to be a widely used backup mechanism to the better discovery methods that occur out-of-band. However, FILS and unsolicited probe responses will most likely only be used in corner cases.




FIGURE 6-7: PSC and 80 MHz channels.



REMEMBER

Out-of-band discovery should be the primary method used, but try and think of the in-band discovery methods as a backup.

Everything Is Bigger in 6 GHz

Just like it's a common truism that “everything is bigger in Texas,” when discussing Wi-Fi, “everything is bigger in 6 GHz.” Especially when talking about bigger and wider channels.

Actually, large Wi-Fi channels are not a new concept. 802.11n (Wi-Fi 4) introduced the use of 40 MHz channels in the 5 GHz frequency band. 802.11ac (Wi-Fi 5) introduced 80 MHz and even 160 MHz channels in the 5 GHz band. Wi-Fi 6 radios that operate in the 5 GHz band can also use these wide channels. But the truth is that there simply is not enough frequency space in 5 GHz for an effective 80 MHz channel reuse plan. As a matter of fact, 20 MHz channel reuse plans are still dominant in 5 GHz instead of 40 MHz.

40 MHz channels are created by bonding together two 20 MHz channels. Channel bonding effectively doubles the frequency bandwidth, meaning double the data rates. And doubling the data rates means double the throughput. Although 20 MHz channel reuse is the norm, 40 MHz channel reuse patterns can be effective in the 5 GHz band with good design best practices: Use all the dynamic frequency selection (DFS) channels with low power and thick walls with high attenuation. On a side note, channel bonding should never be used in 2.4 GHz in the enterprise.

Now, very large 80 and 160 MHz channels are created by bonding together multiple 20 MHz channels. Even though 80 MHz and 160 MHz channels are available in 5 GHz, they should not be used in the enterprise. Bottom line, 80 MHz and 160 MHz channel

deployments do not scale in 5 GHz. There is simply not enough 5 GHz frequency space, and co-channel interference is guaranteed to occur. Performance will drop significantly if 80 MHz channels are deployed on multiple 5 GHz APs in the enterprise.

But remember, “everything is bigger in 6 GHz,” so the potential of large channel reuse plans for 6 GHz is quite astounding. Because of all the available frequency space in 6 GHz, it is expected that 40 MHz and even 80 MHz channel reuse plans will become standard practice. In the United States, 29 new 40 MHz channels can be used in a channel reuse plan. If all 29 of the 40 MHz channels are used, performance degradation due to co-channel interference should be a non-issue. In Europe, there will be 12 new 40 MHz channels available for a reuse plan.

And for the first time, the use of 80 MHz channels will be a reality for enterprise deployments. In the United States, 14 channels are available for an 80 MHz channel reuse plan in 6 GHz. Co-channel interference can probably be limited if all 14 of the 80 MHz channels are used together with careful planning and design. As a matter of fact, I expect that most enterprise vendor APs will have 80 MHz enabled as the default channel size.

You might ask, “What about the noise floor?” Another problem with channel bonding is that it also results in a higher noise floor. The noise floor rises 3 dB when 40 MHz channels are used. If the noise floor is 3 dB higher, then the *signal-to-noise ratio (SNR)* is 3 dB lower, which means that the Wi-Fi radios may shift down to lower modulation data rates. In many cases, this offsets some of the bandwidth gains that the extra frequency space provides. If 80 MHz channels are deployed, the noise floor is 6 dB higher, and the SNR is 6 dB lower. A 6 dB hit on SNR is significant. Would this problem still apply to 6 GHz? Yes, but there is a solution. And it doesn’t involve contacting the authorities about all that loud music next door. Instead, the FCC has defined new transmit power rules that actually favor the use of large 80 MHz channels.

An intriguing difference in 6 GHz will be the new *power spectral density (PSD)* rules that should offset rises in the noise floor caused by channel bonding. PSD is the measure of signal-strength (energy) variations as a function of frequency. A unit of PSD is energy per frequency (width), for example, 5 dBm/MHz.

For 6 GHz channels, the FCC limits radios and antennas by power spectral density (PSD). The FCC will allow a maximum radiated power spectral density of 5 dBm per 1 megahertz for low-power indoor (LPI) APs. Under the control of the LPI APs, clients are permitted a maximum radiated power spectral density of -1 dBm per 1 megahertz. Confused yet?

Let me try and make this a little easier to understand. Signal-to-noise ratio (SNR) is a great metric for a quality RF signal. Figure 6-8 shows that the SNR is simply the difference in decibels between the received signal and the background noise (noise floor) measured in decibels (dBs). If a Wi-Fi radio receives a signal of -70 dBm and the noise floor is measured at -95 dBm, the difference between the received signal and the background noise is 25 dB. Therefore, the SNR is 25 dB.




FIGURE 6-8: Signal-to-noise ratio.

An SNR of 25 dB or greater is considered good signal quality, and an SNR of 10 dB or lower is considered very poor signal quality. An SNR of below 10 dB will likely result in data corruption and retransmission rates as high as 50 percent. To ensure that frames are not corrupted due to a low SNR, most WLAN vendors recommend a minimum SNR of 20 dB for data WLANs and a minimum SNR of 25 dB for WLANs that require voice-grade communications.

In the past, EIRP has always been a constant and the SNR a variable. Every time you double the channel size, the 3 dB rise in the noise floor lowers the SNR. However, the new FCC indoor rules in 6 GHz instead keep PSD constant. And every time you double the channel size, the EIRP is also increased. As shown in Table 6-2, the noise floor still rises 3 dB, but so does the EIRP by 3 dB. The result is an Effective EIRP that stays at the same level, and even more important is that the SNR remains constant. Please note that SNR is dependent on the RF environment, and the table assumes that we have a quality SNR of 25 dB.

TABLE 6-2 EIRP, PSD, and Channel Size for Low-Power Indoor (LPI) APs

EIRP	PSD	Channel Width	Noise Floor	Effective EIRP	Assumed SNR
18 dBm	5 dBm/MHz	20 MHz		18 dBm	25 dB
21 dBm	5 dBm/MHz	40 MHz	+3 dBm	18 dBm	25 dB
24 dBm	5 dBm/MHz	80 MHz	+6 dBm	18 dBm	25 dB
27 dBm	5 dBm/MHz	160 MHz	+9 dBm	18 dBm	25 dB
30 dBm	5 dBm/MHz	320 MHz	+12 dBm	18 dBm	25 dB

I once promised myself that I would not put any formulas in a Dummies booklet, but I will break that promise. The conversion between PSD and EIRP can be calculated with this simple logarithmic formula:

$$\text{EIRP} = \text{PSD (dBm/MHz)} + 10\log(\text{channel width in MHz})$$

The bottom line is that the 5 dBm/MHz rules will compensate for the 3 dB rise in the noise floor when 6 GHz channels are bonded. As a result, 80 MHz channel reuse patterns in the enterprise could become common in countries with the entire 1,200 MHz of 6 GHz frequency space available. In Europe, 40 MHz will probably be more common because there are twelve 40 MHz channels available for a reuse plan, but only six 80 MHz channels. Keep in mind that 6 GHz Wi-Fi has not been field-tested yet; therefore, it remains to be seen if 80 MHz channel reuse patterns in the enterprise become prevalent. But certainly, 40 MHz channel plans will be feasible. So, will 160 MHz channels be used in the enterprise? Probably not at any scale; however, APs deployed in isolated areas could use a 160 MHz channel.

Another thing to consider when deploying big channels in 6 GHz is the selection of the primary channels. When channel bonding is used, one 20 MHz channel is selected as the *primary channel* while all the other channels are known as secondary channels. The primary channel is used for the transmission of management and control frames. On the other hand, data frames are modulated across the entire bonded channel. In 5 GHz, almost always, the first 20 MHz of any 40, 80, or 160 MHz channel is the primary.

This will be different in 6 GHz; instead, the second 20 MHz of any 40, 80, or 160 MHz channel will usually be the primary. Take a look at Figure 6-9, which depicts the U-NII-5 band in 6 GHz. I refer you to the earlier discussion about preferred scanning channels (PSCs) used for active discovery.




FIGURE 6-9: Preferred scanning channels and primary channels.

If PSC is used, the PSC channels must also serve as the primary channels when channel bonding is used for 40, 80, or 160 MHz channels. Please note that the positioning of the PSC channels is the *second* 20 MHz of a bonded channel. If enabled, the PSC channels also function as the primary channel for half of the 40 MHz channels and all 80 MHz channels, as shown in Figure 6-8.



Because PSC is expected to be enabled on most APs, my recommendation would be to always use the second 20 MHz channel as the primary in any 40, 80, or 160 MHz channels.

- » Understanding WPA3-Personal
- » Achieving WPA3-Enterprise
- » Reviewing Enhanced Open
- » Outlining some key 6 GHz security takeaways
- » Implementing WIPS

Chapter 7

Security in a 6 GHz Wi-Fi 6E World

Prior to the expected 6 GHz Wi-Fi spectrum bonanza, ongoing enhancements were made toward shoring up Wi-Fi security with both WPA3 and Enhanced Open for all Wi-Fi frequencies. As is to be expected, there will be Wi-Fi security considerations when deploying Wi-Fi in the 6 GHz frequency band. The Wi-Fi Alliance will require WPA3 security certification for Wi-Fi 6E devices that will operate in the 6 GHz band. Furthermore, support for the Enhanced Open security certification will also be mandatory.

WPA3-Personal

In August 2019, the Wi-Fi Alliance began testing access points (APs) and clients for the Wi-Fi Certified WPA3 certification. Wi-Fi Protected Access 3 (WPA3) defines enhancements to the existing WPA2 security capabilities for 802.11 radios.

By far, the most significant change defined by WPA3 is the replacement of preshared key (PSK) authentication with *Simultaneous Authentication of Equals* (SAE), which is resistant to offline dictionary attacks. SAE is based on a Dragonfly key exchange.

Dragonfly is a patent-free and royalty-free technology that uses a zero-knowledge proof key exchange, which means a user or device must prove knowledge of a password without revealing the password. Think of SAE as a more secure PSK authentication method. The goal is to provide the same user experience by still using a passphrase. However, the SAE protocol exchange protects the passphrase from brute-force dictionary attacks. The passphrase is never sent between Wi-Fi devices during the SAE exchange.

As shown in Figure 7-1, an SAE process consists of a commitment message exchange and a confirmation message exchange. The commitment exchange is used to force each radio to commit to a single guess of the passphrase. Next, the confirmation exchange is used to prove that the password guess was correct. The passphrase is used in SAE to deterministically compute a secret password element used for the authentication and key exchange protocol. After the SAE exchanges are complete, a unique *pairwise master key (PMK)* is derived and installed on both the AP and the client station. The PMK is the seeding material for the 4-Way Handshake that is used to generate dynamic encryption keys. SAE authentication is performed prior to association. Once the PMK is created and the association process completes, the AP and the client can then commence a 4-Way Handshake to create a *pairwise transient key (PTK)*. The PTK is the dynamically generated key used to encrypt unicast traffic.

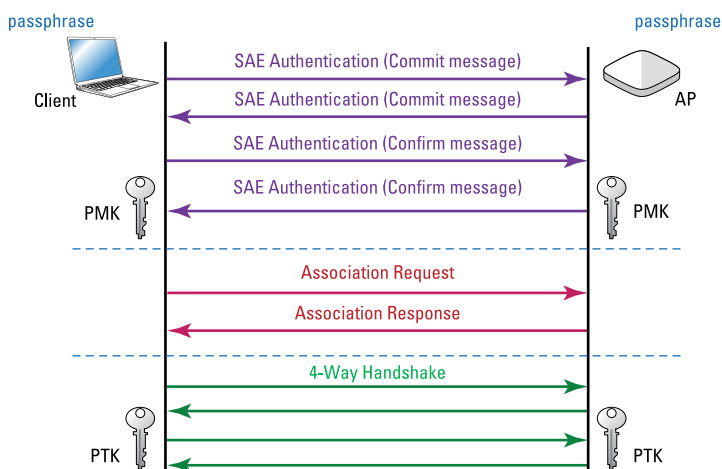


FIGURE 7-1: Simultaneous Authentication of Equals.

You're basically asking the devices to shout the first thing that comes to their heads and hope it ends in a technologically romantic meet-cute. Minus swelling background music and the end credits, of course.

WPA3-Personal enhances Wi-Fi security for home users and environments where 802.1X is not an option. From the perspective of the user, the connection experience remains the same. A passphrase is still used to connect; however, the SAE protocol exchange protects the passphrase from brute-force dictionary attacks. WPA3-Personal defines two modes of operation:

- » **WPA3-Personal Only:** This mode completely replaces WPA2 PSK authentication and requires the use of SAE authentication. This mode would only be enabled on the AP if all clients were WPA3-capable. Management frame protection (MFP) is required for both APs and clients operating in this mode.
- » **WPA-3 Personal Transition:** The transitional mode allows for backward compatibility with WPA2-Personal. This mode allows for WPA2-Personal clients to connect to the same SSID as WPA3-Personal clients. The clients use the same passphrase. However, the WPA2 clients connect with PSK authentication, and the WPA3 clients connect with SAE authentication. In this mode, MFP is used by the WPA3 clients but not necessarily by the WPA2 clients.

WPA3-Enterprise

Unlike WPA3-Personal, where an entirely new authentication method has been designated, WPA3-Enterprise still leverages 802.1X/EAP for enterprise-grade authentication. In other words, the enterprise-grade authentication process remains the same. The two main enhancements are support for MFP and an optional enhanced cryptographic mode. WPA3-Enterprise defines three modes of operation:

- » **WPA-3 Enterprise Only:** 802.1X/EAP authentication remains the same. However, this mode would only be enabled on the AP if all clients were WPA3-capable. Management frame protection (MFP) is required for both APs and clients operating in this mode.

- » **WPA-3 Enterprise Transition:** The transitional mode allows for backward compatibility with WPA2-Enterprise. This allows WPA2-Enterprise clients to connect to the same SSID as WPA3-Enterprise clients. 802.1X/EAP authentication remains the same. However, in this mode, MFP is used by the WPA3 clients but not necessarily by the WPA2 clients.
- » **WPA-3 Enterprise 192-bit:** This mode may be deployed in sensitive enterprise environments to further protect Wi-Fi networks with higher security requirements such as government, defense, and industrial. This is an optional mode using 192-bit minimum-strength security protocols and cryptographic tools to better protect sensitive data.

Despite the transitional modes offered by WPA3 currently, tactical deployments of WPA3 security are rare in the enterprise. WPA2-Enterprise still offers almost the same level of 802.1X/EAP authentication security as WPA3-Enterprise. WPA3-Personal and the use of SAE is growing in the consumer market. However, SAE hasn't been embraced yet in the enterprise where PSK authentication may or may not be used. The bulk of the enterprise Wi-Fi client population supports and continues to use WPA2 security. Additionally, even though WPA3 firmware upgrades are possible for older client devices, most client vendors may never offer a WPA3 firmware update for a client device that is three or more years old.



The Wi-Fi Alliance mandates support for WPA3 security for Wi-Fi 6 certification, meaning that all 802.11ax radios must support WPA3. Furthermore, as of July 1, 2020, the Wi-Fi Alliance mandates support of WPA3 security for all future certifications. In other words, all the Wi-Fi radios currently hitting the market must support WPA3. Adoption is still another matter.

Enhanced Open

Traditionally, Wi-Fi hotspots and guest WLANs have used open security without encryption or authentication. Passpoint security is catching on in the Wi-Fi public access marketplace, but we still have a ways to go with that. The Wi-Fi CERTIFIED Enhanced Open certification defines improved data privacy in open Wi-Fi networks. This certification is based on the *Opportunistic Wireless*

Encryption (OWE) protocol. The OWE protocol integrates established cryptography mechanisms to provide each user with unique individual encryption, protecting the data exchange between the user and the AP. With OWE, standard open authentication and association occur, and then the 4-Way Handshake process generates the necessary keys for encryption.

The OWE experience for the user is the same as open security because there is no need to enter a password or passphrase before joining the network. Data privacy is provided, and malicious eavesdropping attacks are mitigated because the 802.11 data frames are encrypted. But please understand that there is zero authentication security. Enhanced Open is not part of WPA3 and is an entirely different and optional security certification for 2.4 GHz and 5 GHz frequency bands.

Understand that Enhanced Open meets only half of the requirements for well-rounded Wi-Fi security. OWE does provide encryption and data privacy, but there is no authentication whatsoever.



REMEMBER

Enhanced Open is an optional security certification. As a result, many WLAN vendors still do not support OWE, and client-side support is marginal at best but growing. Therefore, tactical deployments of OWE in the 2.4 and 5 GHz frequency bands are currently scarce. However, the Enhanced Open certification is mandated for 6 GHz.

Key 6 GHz Security Takeaways

So, as to be expected, there are Wi-Fi security considerations when deploying Wi-Fi in the 6 GHz frequency band. The Wi-Fi Alliance will require WPA3 security certification for Wi-Fi 6E devices that will operate in the 6 GHz band. However, there will be no backward compatibility support for WPA2 security. Furthermore, the Enhanced Open certification will mandate support for Opportunistic Wireless Encryption (OWE) in 6 GHz. The good news is that all 802.11 data traffic will always be encrypted in 6 GHz to provide data privacy. Furthermore, all 802.11 management traffic will be protected by MFP to help mitigate well-known layer 2 denial-of-service (DoS) attacks.

Because OWE support will be mandatory, there will not be any “open” security SSIDs operating in 6 GHz. I personally have never been a big fan of OWE because it only provides encryption but not authentication. WPA3–Personal or WPA3–Enterprise are better options because authentication is also a requirement. The bottom line is that open networks are not permitted in 6 GHz and all data traffic will be encrypted. This will have implications for existing businesses that are currently using open guest access in the legacy bands.

Because there is no backward compatibility for WPA2, there will be no support for PSK authentication. Once again, the WPA3–Personal replacement for PSK is Simultaneous Authentication of Equals (SAE). WPA3–Enterprise will still use 802.1X. Management frame protection (MFP) will also be required. Because there is no backward compatibility for WPA2, there will be no need for either the WPA3–Personal transition mode or the WPA3–Enterprise transition mode.


So, what are the most critical takeaways when implementing Wi-Fi security in the 6 GHz band? I’m glad you asked, here’s a list:

- » **Different layers of Wi-Fi security:** It is my belief that because the existing 15 billion Wi-Fi clients will never be able to connect to 6 GHz, it appears likely that different levels of security will be used on the different frequency bands in the enterprise. WPA3 will indeed be used in 6 GHz. Yet, despite the support for WPA3 transition modes in the legacy bands, WPA2 will likely remain prevalent in the 2.4 GHz and 5 GHz bands for a very long time.
- » **Different SSIDs per frequency:** This means that different SSIDs with different levels of security will be used on the various bands. For example, as depicted in Figure 7-2, an *employee* SSID using WPA2–Enterprise and a *guest* SSID using open security are used for the 2.4 and 5 GHz bands. However, the 6 GHz band requires different SSIDs and security: *employee-6* using WPA3–Enterprise and *guest-6* using Enhanced Open.

I think it will take time, but the anticipated wide adoption of 6 GHz enterprise deployments hopefully will accelerate the transition to WPA3 security in the other frequency bands. In the meantime, I expect various levels of security across the three bands.

5 GHz channel (100)

- SSID: **employee**
 - WPA2-Enterprise
- SSID: **guest**
 - Open



6 GHz channel (37)

- SSID: **employee-6**
 - WPA3-Enterprise
- SSID: **guest-6**
 - Enhanced Open

2.4 GHz channel (1)

- SSID: **employee**
 - WPA2-Enterprise
- SSID: **guest**
 - Open



FIGURE 7-2: Different SSIDs and Security across three frequency bands.

WIPS

The big buzz-phrase in Wi-Fi security has always been the rogue AP: a potential open and unsecured gateway straight into the entire network that the company wants to protect. Wireless can be an intrusive technology, and if wired data ports at a business are not controlled, any individual (including an employee) can install a rogue AP. A rogue AP is any unauthorized Wi-Fi device that is not under the management of the proper network administrators.

The same skull and crossbones symbol that is used by Caribbean pirates is often also used as an icon in WIPS solutions to represent a rogue AP. Using a ship to commit acts of robbery and violence against a coastal area or other ships is the definition of piracy. The ships used to commit these acts are pirate ships. Using a wireless rogue device for data theft, data destruction, loss of services, and other attacks are all acts of wireless piracy. A rogue AP is effectively a pirate ship, albeit with fewer parrots and barrels of rum. As shown in Figure 7-3, the advent of 6 GHz devices brings a new breed of pirate ships that pose a hostile threat.



REMEMBER

In today's world, a *wireless intrusion prevention system (WIPS)* is necessary to monitor for rogue devices and other wireless attacks. 6 GHz Wi-Fi presents many new challenges for WIPS security, including:

- » **Sensors:** The sensors deployed with current WIPS solutions use 2.4 and 5 GHz radios. In other words, they do not have 6 GHz radios in the existing sensors to monitor for attacks. Vendors that offer Wi-Fi 6E APs with tri-frequency sensor scanning capabilities will take the lead in 6 GHz WIPS security.




FIGURE 7-3: Pirates of the 6 GHz Caribbean.

» **Scanning:** There are now 59 channels in 6 GHz that must be monitored for potential attacks. The scanning intervals will be much longer. WIPS vendors will most likely offer multiple options of scanning intervals for both full-time dedicated 6 GHz sensors and APs that function as part-time sensors when using off-channel scanning.

» **Mitigation:** Rogue mitigation (also known as rogue containment) is accomplished wirelessly when WIPS' sensors become active and transmit layer 2 deauthentication frames that spoof the MAC addresses of the rogue APs and rogue clients. WIPS solutions use a known layer 2 denial-of-service attack as a countermeasure. However, WPA3 is mandated in 6 GHz, and management frame protection (MFP) is required. MFP will render many of the common layer 2 mitigation techniques useless. WIPS vendors will need to rely on alternative methods of disabling rogue APs such as switch port suppression.



TIP

Would you like to learn more about WIPS security? May I suggest the eBook, *Wireless Intrusion Prevention Systems (WIPS) For Dummies*, written by yours truly?

- » Understanding Wi-Fi 6 client capabilities
- » Achieving MultiGig wired speeds
- » Providing Power over Ethernet
- » Debating 4×4:4 versus 8×8:8
- » Measuring 6 GHz distance

Chapter 8

Wi-Fi 6 and 6E Key Questions

New standards, new questions. This quandary is nothing new for tech experts and new parents alike. We can at least help the tech experts here. In this chapter, you learn about many of the most commonly asked real-world deployment questions regarding Wi-Fi 6 and Wi-Fi 6E.

Clients

The first client question that I am being asked is, “When will we see Wi-Fi 6E clients (6 GHz), and how fast will we see them in the enterprise?” Believe it or not, Wi-Fi 6 clients have been in the marketplace for two years. And now, Wi-Fi 6E radios with 6 GHz functionality are finding their way into smartphones, tablets, and laptops. Samsung released the Galaxy S21 Ultra, the first Wi-Fi 6E smartphone, into the market in January of 2021. At least seven more smartphones are expected to be available between then and February of 2022. Laptop manufacturers such as Dell and Lenovo already offer new models with Wi-Fi 6E radios.

The introduction of new technology is often consumer-driven. While the enterprise is often slow at updating Wi-Fi clients, employees often force the issue because they bring their devices into the workplace. The catchphrase of *bring your own device* (BYOD) is a direct result of employee expectations connecting to the corporate WLAN with multiple personal mobile devices. Don't forget that we now live in a world of an infinite enterprise where many employees are currently working at home or in a hybrid home-and-office schedule that requires both locations to handle the most modern technology. Therefore, much of the client-side implementation of Wi-Fi 6E clients will be consumer-driven.



REMEMBER

Industry analysts all agree that the Wi-Fi 6 and 6E technology growth will be fast and furious (and probably as well scripted as their namesake movies). For example, in 2021, over 50 percent of all Wi-Fi shipments are expected to be Wi-Fi 6, and more than 338 million Wi-Fi 6E devices will enter the market. Technology research group, IDC, predicts 5.2 billion Wi-Fi 6 product shipments by 2025, 41% of which will be Wi-Fi 6E devices.

Surprisingly, I am also often asked, “Will there be any 8×8:8 Wi-Fi 6 or 6E clients?” Most Wi-Fi mobile client devices such as smartphones will use 2×2:2 radios because an 8×8:8 radio would drain battery life. The majority of Wi-Fi client devices are 2×2:2. In the future, you might see some 4×4:4 client radios in high-end laptops.

The other big question that I get all the time is, “Will there be any performance benefit for legacy clients when Wi-Fi 6 APs (access points) are deployed?” The answer is yes and no. First, the bad news: Legacy 802.11n/ac clients do not support Wi-Fi 6 mechanisms such as OFDMA. Therefore, the legacy clients will continue to use single-user communications when connected to a Wi-Fi 6 AP. Wi-Fi 6 clients are needed to take full advantage of 802.11ax high efficiency capabilities such as multi-user OFDMA. However, there is still good news for the legacy clients for three reasons:

- » **AP hardware:** Although there are no physical (PHY) layer improvements with legacy clients, you will see performance improvements as a result of newer hardware capabilities of the new Wi-Fi 6 APs, such as stronger CPUs and better memory handling.

- » **Airtime availability:** As more Wi-Fi 6 clients are mixed into the 2.4 and 5 GHz client population, the efficiency improvements gained by Wi-Fi 6 client devices will free valuable airtime for the legacy clients, therefore improving the overall efficiency of the wireless network.
- » **Mesh:** Remember, legacy 802.11a/b/g/n/ac clients are not supported in 6 GHz. However, they may indirectly benefit from having their traffic carried in a 6 GHz mesh backhaul link. Most Wi-Fi 6 enterprise APs are expected to be tri-frequency.

On a related note, I am also often asked, “Is it possible to upgrade Wi-Fi 6 devices that operate in the 2.4 and 5 GHz frequency bands to also transmit in the 6 GHz frequency band?” The simple answer is no because a hardware upgrade will be required.

MultiGig

“Will MultiGig Ethernet be a necessity?” With each new generation of Wi-Fi technology and higher data rates, we’ve seen various bandwidth claims made regarding the wired uplink connection between an AP and an access switch. Because Wi-Fi data rates have risen dramatically, the worrisome claims are that a standard 1 Gbps (gigabits per second) wired uplink will become a bottleneck. Take a look at this historically. As early as 2009, when 802.11n debuted, vendors claimed we would need to aggregate GbE ports with two cables to prevent bottlenecks. Later, when 802.11ac debuted, many enterprise switch vendors claimed that everyone should upgrade their switches to support 2.5 GbE uplinks. All of these past predictions of access-layer bottlenecks have not happened.

Prior to Wi-Fi 6 (802.11ax), the only time a 1 Gbps uplink has not been sufficient is in laboratory test environments or very unique corner cases. Bandwidth bottlenecks almost *never* occur at the access layer. However, bandwidth bottlenecks can certainly occur on the wired network due to poor wired network design. The number one bandwidth bottleneck is usually the WAN uplink at any remote site. But it is safe to say that the Wi-Fi will always be blamed first despite the inadequate WAN bandwidth.



As shown in Figure 8-1, the 802.3bz standard (also known as MultiGig Ethernet) defines bandwidth capabilities of up to 2.5 Gbps and 5 Gbps over Cat5e and Cat6 copper cables.




FIGURE 8-1: MultiGig Ethernet – 802.3bz.



Note that 10 Gbps bandwidth is even possible but requires network cable upgrades to Cat6a or Cat7. Enterprise vendors are now actively pushing sales of access switches that support these MultiGig speeds.

So, one question is, “Will we need 2.5 Gbps Ethernet ports for Wi-Fi 6 APs?” The whole point of Wi-Fi 6 (802.11ax) is better spectrum efficiency and a reduction in airtime consumption. Logic dictates that if Wi-Fi becomes more efficient, the user traffic generated by a dual-frequency Wi-Fi 6 AP could potentially exceed 1 Gbps. A more likely cause of a potential bottleneck is that tri-frequency band Wi-Fi 6E APs have arrived. Will the wireless traffic coming from three radios surpass a 1 Gbps uplink? The expected use of wide channels in 6 GHz means higher throughput potential, as well. The fear is that a standard Gigabit Ethernet wired uplink port could be a bottleneck, and therefore 2.5 Gbps uplink ports will be needed. As a precaution, enterprise Wi-Fi 6 and 6E APs now include at least one 802.3bz MultiGig Ethernet port capable of a 2.5 or 5 Gbps wired uplink. Think of this step as future-proofing (or at least future-delaying).

In the real world, we probably will still not exceed 1 Gbps for some time because of these three reasons:

- » **Wi-Fi 6 and 6E client population:** Even though the chipset vendors are aggressively making Wi-Fi 6 and 6E client radios available, we expect some time to pass before the bulk of the enterprise client population is dominated by Wi-Fi 6 and 6E clients.
- » **Legacy clients drag us down:** 802.11ax in the 2.4 and 5 GHz bands requires backward compatibility with 802.11/a/b/g/n/ac, which means that RTS/CTS protection mechanisms must be used. RTS/CTS creates overhead and consumes airtime.
- » **Medium contention:** Under controlled conditions, where only one active client on each of the three frequency bands is transferring data on a tri-band Wi-Fi 6E AP, a single 1 Gbps Ethernet uplink is indeed insufficient. But as multiple clients participate and contend for the medium on all three frequencies, the conditions are rare for 1 Gbps Ethernet uplink to become a bottleneck.

So, will customers have to upgrade their switches to have Multi-Gig capabilities? As I have already discussed, past gloom and doom predictions of access-layer bottlenecks have not come true. Although historically 1 Gbps uplinks have been more than enough, I will predict that 2.5 Gbps uplinks may be needed eventually. In the future, as Wi-Fi 6 client populations grow and as WLAN vendors add tri-band radios into their APs, 1 Gbps uplinks may no longer be sufficient. Once again, think of 2.5 Gbps MultiGig Ethernet as future-proofing. By the way, any vendor claims that we need 10 Gbps uplinks for Wi-Fi 6E is a marketing fantasy.

Power over Ethernet

Probably the much more important conversation about the relationship between switches and Wi-Fi 6 and 6E APs is Power over Ethernet (PoE) requirements. “Will Wi-Fi 6 APs work with standard 802.3af PoE?” Enterprise Wi-Fi manufacturers have added more radio chains to their Wi-Fi 6 APs. Many of the current Wi-Fi 6 APs are dual-band 4×4:4 APs, and there are even 8×8:8 APs. Wi-Fi 6 APs also require much more processing power than previous generations of enterprise APs. The extra radio chains and processor capabilities require more power. The 15.4 watts (W) provided per port by standard 802.11af PoE is not adequate for 4×4:4 APs, and therefore 802.3at (PoE Plus) power is necessary. PoE Plus-capable

switches can provide up to 30 watts of power per Ethernet port. PoE Plus-powered ports for 4×4:4 APs should be considered a standard requirement.

The exception to this requirement is dual-band 2×2:2 802.11ax APs. In most cases, standard PoE of 15.4 watts will be sufficient to power these APs. However, Wi-Fi 6E introduces tri-frequency capabilities. So, 15.4 watts (W) is not adequate for three 2×2:2 radios housed in the same AP. Therefore, 802.3at (PoE Plus) power is necessary for pretty much all Wi-Fi 6E AP form factors.



TIP

If a business does not have PoE Plus-capable switches, they will have to upgrade if they deploy any dual-band 4×4:4 APs or tri-radio Wi-Fi 6E APs. There is a good chance that many enterprise businesses already have access switches with PoE Plus capabilities. But do the switches have a total power budget for a 1:1 replacement for all the APs that require PoE Plus? I am worried that many businesses will suddenly be exceeding the overall power budgets of the switches. Enterprise Wi-Fi vendors commonly receive support calls from customers complaining that all of a sudden, APs randomly begin to reboot. In most cases, the root cause of random rebooting of APs is that the switch power budget has been eclipsed. Very often, if an AP cannot get the power that it needs, the AP will reboot and try again. The power budget of a switch or multiple switches should be monitored to make sure that all devices can maintain power. Active power budget information can usually be seen from the command line of a switch or the GUI interface or monitored by a cloud management solution such as ExtremeCloud IQ. An upgrade to most Wi-Fi 6 and 6E APs will at the very least require a recalculation of PoE power budgets. As WLAN vendors add more radio chains, dual-band radios, and now tri-band radios, PoE power budget management will be of even greater importance moving forward.

Some WLAN vendors sell 8×8:8 Wi-Fi 6 APs, and the PoE power requirements are even more substantial. In some cases, these 8×8:8 APs require 31 watts of power or more, which means that even PoE Plus power will not be sufficient. While some of the 8×8:8 APs can be powered by 802.11at (PoE Plus), there is often some sort of downgrade functionality. For example, you might experience the loss of USB availability, BLE radio, or some other functionality. It is also conceivable that when higher-end Wi-Fi 6E APs hit the market in 2022, that as much as 45 watts will be needed to power 4×4:4 tri-frequency APs. In that case, switches that support 802.3bt PoE will be a requirement.



802.3bt is a new PoE standard that defines per-port power-source capabilities of 45 watts, 60 watts, 75 watts, and even 90 watts. Enterprise manufacturers now sell switches that support 802.3bt. Keep in mind that this technology is quite expensive. For now, PoE Plus (802.3at) is sufficient to power the majority of Wi-Fi 6 and 6E APs. In the future, as WLAN vendors manufacture APs with even more radios under the hood, the need for 802.3bt power might become relevant.

4x4:4 versus 8x8:8

Probably a question that seems to always pop-up is, “Which is better? 8x8:8 APs or 4x4:4 APs?” Several WLAN vendors are selling Wi-Fi 6 APs that have a 2.4 GHz radio that is 4x4:4 and a 5 GHz radio that is 8x8:8. These WLAN vendors are, of course, forging marketing pitches for 8x8:8. Eight must be better than four — right? More is always better!

Theoretically, an 8x8:8 AP can modulate data on all eight radio chains to a single client resulting in some substantially high data rates. The problem is that there will never be any 8x8:8 mobile client devices due to the drain on battery life. As shown in Figure 8-2, we live in a world where the bulk of Wi-Fi client devices are 2x2:2. And depending on existing conditions, the clients will often downgrade to 1x1:1 communication.



FIGURE 8-2: 2x2:2 client capabilities and operational functionality.

So, the primary advantage of an 8x8:8 AP over a 4x4:4 AP is MU-MIMO functionality. An 8x8:8 AP could modulate two independent data streams to each of four 802.11ax 2x2:2 clients that support downlink MU-MIMO. Also, an 8x8:8 AP could transmit downlink one unique modulated stream of data to each of eight 802.11ax

clients, simultaneously. While this sounds good in theory, I will refer you to our MU-MIMO discussion in Chapter 3. MU-MIMO requires spatial diversity. Even though all Wi-Fi 6 clients will support downlink MU-MIMO, most modern-day enterprise deployments of Wi-Fi involve a high density of users and devices that is not conducive for MU-MIMO conditions. MU-MIMO requires a sizable physical distance between the clients, as well as the AP, for spatial diversity. Check out Chapter 3 for more information.



TIP

You might read some marketing hype that 8×8:8 APs will support more clients than a 4×4:4 AP. As you learned in Chapter 2, OFDMA is the multi-user technology that holds the most promise for efficiency improvements. Regardless of the number of radio chains and stream count, all Wi-Fi 6 APs will support the same number of OFDMA clients during a transmission opportunity (TXOP).

8×8:8 APs are very expensive and are a huge drain on the PoE power budget. While the MU-MIMO gains sound enticing in theory, the reality is that in most indoor enterprise deployments, an 8×8:8 AP offers no real practical advantage over a less expensive 4×4:4 Wi-Fi 6 AP.

And guess what? All the major enterprise WLAN vendors sell dual-band (Wi-Fi 6) 2×2:2 APs and now tri-band (Wi-Fi 6E) 2×2:2 APs. The 2×2:2 APs are actually quite popular because of the lower cost. And 2×2:2 Wi-Fi 6 and 6E APs still offer the bulk of the 802.11ax benefits, including OFDMA, TWT, 1024-QAM, and more. Regardless of the number of radio chains and regardless of stream count, all Wi-Fi 6 and 6E APs support the same number of OFDMA clients during a transmission opportunity (TXOP).

6 GHz Range

The 2.4 GHz band will still be considered a “best effort” frequency band, and the 5 GHz channels will be used for clients that require higher performance metrics. However, the potential of 6 GHz is quite astounding due to all the newly available frequency space (1,200 MHz). As 6 GHz-capable client populations grow, WLANs will also have to be designed for indoor 6 GHz coverage. Vendors are already manufacturing APs with radios for all three frequencies. And therefore, a valid question specific to Wi-Fi 6E is, “Will I have to redesign my network because 6 GHz will not have the same coverage range?”

Because of the laws of physics, an electromagnetic signal will attenuate as it travels, despite the lack of attenuation caused by obstructions, absorption, reflection, diffraction, and so on. *Free space path loss (FSPL)* is the loss of signal strength caused by the natural broadening of the waves, often referred to as beam divergence. RF signal energy spreads over larger areas as the signal travels farther away from an antenna, and as a result, the strength of the signal attenuates. One way to illustrate free space path loss is to use a balloon analogy. Before a balloon is filled with air, it remains small but has a dense rubber thickness. After the balloon is inflated and has grown and spread in size, the rubber becomes very thin. RF signals lose strength in much the same manner. And due to FSPL, an RF signal loses the most power in the first meter it travels.

A decibel (dB) is a logarithmic measure of signal strength gain or loss. Figure 8-3 shows that a 2.4 GHz signal loses about 40 dB in the first meter. The reason that the *effective range* of a 5 GHz AP is much smaller is that the 5 GHz signal attenuates 47 dB in the same first meter. An easier way to explain the difference is that the 5 GHz signal attenuates five times more than a 2.4 GHz signal in the first meter. Luckily, this loss in signal strength is logarithmic and not linear; thus, the amplitude does not decrease as much in a second segment of equal length as it decreases in the first segment. There is a fancy logarithmic equation to calculate free space path loss; however, the *6 dB rule* is an easy way to estimate FSPL. The 6 dB rule states that doubling the distance will result in a loss of amplitude of 6 dB, regardless of the frequency. Therefore, at 2 meters, the path loss is 46 dB for 2.4 GHz, 53 dB for 5 GHz, and 55 dB for 6 GHz.




FIGURE 8-3: Free space path loss (FSPL) in the first meter.

In the past, coverage planning has been for two bands. For dual-frequency APs, the planning and validation of -65 dBm or -70 dBm coverage was based on the 5 GHz radio. The reason is that the effective range of 5 GHz is much smaller than 2.4 GHz. Therefore, using the lowest common denominator of 5 GHz was preferred when planning for coverage. The good news is that the effective range difference between 6 GHz and 5 GHz is not as significant as the difference between 5 GHz and 2.4 GHz.

On average, a 6 GHz signal attenuates about 2 dB more than a 5 GHz signal in the first meter. Of course, the 6 GHz band is big, so it does depend on what channel is being used. For example, as depicted in Figure 8-4, in the UNII-5 band of 6 GHz, the path loss in the first meter is about 48 dB, which is only a single dB difference from the average 5 GHz path loss. The first meter path loss in the center of the 6 GHz band is closer to 49 dB.




FIGURE 8-4: Free space path loss (FSPL) in the first meter – 6 GHz.

Whatever the frequency or channel, remember, after the initial first meter, the 6 dB rule states that doubling the distance will result in a loss of amplitude of 6 dB regardless of the frequency.

The bottom line is that the effective range difference between 6 GHz and 5 GHz is not going to be a serious concern in most indoor Wi-Fi deployments. So, will you have to overhaul your Wi-Fi network due to 6 GHz coverage concerns? In most cases, probably not. However, I am a big proponent of proper WLAN planning and design no matter what the frequency. The bulk of troubleshooting calls can be prevented if a WLAN is well planned and designed before deployment. Just as important is a post-deployment validation survey to verify the WLAN design.

Needless to say, it is safe to assume that all the various Wi-Fi predictive modeling solutions such as Ekahau Survey, iBWave Design, and TamoGraph Site Survey will offer 6 GHz design capabilities. For Greenfield deployments with tri-band APs that include 6 GHz radios, the lowest common denominator for coverage design will now be 6 GHz.

Chapter 9

Ten Things to Know about Wi-Fi 6 and 6E

This book shared a huge amount of information about Wi-Fi 6 and its brand-new technological benefits. Science states that this book can theoretically contain only ten more topics before exceeding known knowledge limits. Therefore, we hit the limit here when I talk about ten key things to keep in mind about Wi-Fi 6 and Wi-Fi 6E.

- » **It is a Wi-Fi paradigm shift.** Wi-Fi 6 does not just push the envelope with regard to Wi-Fi speeds — up to 10 Gbps. This new generation introduces numerous performance improvements as well. In fact, it has been dubbed “High Efficiency” unlike previous versions that were labeled “High Throughput.” Wi-Fi 6 is not just about better throughput. Wi-Fi 6 substantially improves capacity, provides better coverage, and reduces congestion in Wi-Fi networks.
- » **Is it backward or forward compatible?** Unlike the 802.11ac standard, 802.11ax and Wi-Fi 6 support both 2.4 and 5 GHz wireless devices, so 802.11n (and potentially 802.11g and 802.11b) devices can run on Wi-Fi 6 networks. This factor is critical for many legacy specialized devices, particularly in healthcare and manufacturing verticals that tend to move

slowly to update their devices. On the other hand, Wi-Fi 6E is the first generation to offer Wi-Fi capabilities in the 6 GHz frequency band. A key difference of using the 6 GHz frequency band for 802.11ax technology is there is no need for backward compatibility. The 6 GHz frequency band will be a “pure” 802.11ax band for Wi-Fi communication. Legacy 802.11a/b/g/n/ac clients are not supported in 6 GHz.

- » **Do not confuse your multi-user technologies.** The term multi-user (MU) simply means that transmissions between an access point (AP) and multiple clients can occur at the same time, depending on the supported technology. However, the MU terminology can be very confusing when discussing Wi-Fi 6. Multi-user capabilities exist for both OFDMA and MU-MIMO. These differences are key between both Wi-Fi 6 multi-user technologies.
- » **RU ready for simultaneous multi-user Wi-Fi access?** Multi-user *orthogonal frequency division multiple access (OFDMA)* is easily the most important new capability introduced with Wi-Fi 6. It subdivides a channel into smaller frequency allocations, called *resource units (RUs)*, thereby enabling an AP to synchronize communication (uplink and downlink) with multiple individual clients assigned to the RUs.
- » **MU-MIMO enhancements are here.** More MU-MIMO clients can communicate with an AP at the same time. Wi-Fi 5 is limited to a MU-MIMO group of only four clients, whereas Wi-Fi 6 potentially supports up to 8×8:8 MU-MIMO in both downlink and uplink, which allows it to serve up to eight users simultaneously. MU-MIMO is a great multi-user technology for PtMP wireless bridge links between buildings.
- » **No qualms about higher data speeds.** Wi-Fi 6 supports 1024-QAM modulation and coding schemes (MCS) that define higher data rates providing a potential 20 percent increase in data throughput over 256-QAM (introduced in Wi-Fi 5). Additionally, because of all the available frequency space in 6 GHz, Wi-Fi 6E provides new opportunities for high-bandwidth wireless applications. The commonplace use of 80 MHz channels in the enterprise (and sometimes 160 MHz channels) bring us faster data speeds in the enterprise.

- » **You need more power.** Many Wi-Fi 6 APs are dual-band 4x4:4 APs, and some are even 8x8:8 APs. The extra radio chains and processor capabilities require more power. The 15.4 watts provided by standard 802.11af PoE is not adequate. 802.3at (PoE+) power is usually required. PoE+ power for 4x4:4 APs dual-frequency APs should be considered a standard requirement. Additionally, 2x2:2 Wi-Fi 6E APs with tri-frequency radios require PoE+. Upgrades of access-layer switches may be necessary and the recalculation of PoE power budgets is highly recommended.
- » **Wi-Fi 6E is a spectrum bonanza.** The new 1,200 MHz of spectrum provided by Wi-Fi 6E offers an enhanced user experience in all enterprise verticals, including K-12 and higher education, retail, manufacturing, and healthcare. The potential of 1,200 MHz of new frequency space for Wi-Fi communications is mind-boggling. Opening the 6 GHz frequency space for Wi-Fi communication is expected to bring hundreds of billions of dollars of economic value to the worldwide economy.
- » **Wi-Fi 6E clients need to discover 6 GHz APs.** Because there are so many channels in the 6 GHz band, client probing can take a considerable amount of time. A tri-band AP can inform a Wi-Fi 6E client actively scanning the 2.4 GHz or 5 GHz band about existing 6 GHz radios that are co-located in the AP using *Reduced Neighbor Reports* (RNR) and Multi-BSSID beacons. Wi-Fi 6 also offers three potential methods for in-band discovery of 6 GHz APs as a backup to the out-of-band discovery methods.
- » **You may need two levels of Wi-Fi security.** The Wi-Fi Alliance will require WPA3 security certification for Wi-Fi 6E devices that will operate in the 6 GHz band. Furthermore, support for Opportunistic Wireless Encryption (OWE) and the Enhanced Open security certification will also be mandatory in 6 GHz. Because there is no backward compatibility for WPA2 in 6 GHz, there will be no support for PSK authentication. The WPA3-Personal replacement for PSK is Simultaneous Authentication of Equals (SAE). It looks like different levels of security will be used on the different frequency bands in the enterprise. WPA3 is indeed used in 6 GHz. Yet, despite the support for WPA3 transition modes in the legacy bands, WPA2 will likely remain prevalent in the 2.4 GHz and 5 GHz bands for a very long time.

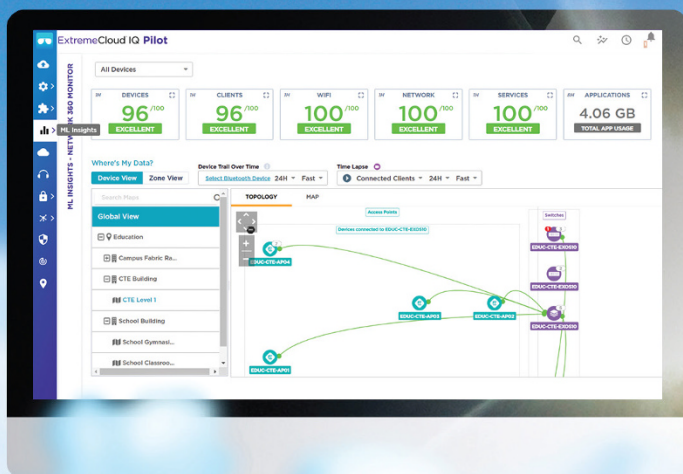


WELCOME TO ExtremeCloud™ IQ

For network administrators looking for unified management of access points, switches, and routers, ExtremeCloud IQ is a cloud-driven network management application that:

- Simplifies network operations and troubleshooting through an easy to use and intuitive interface, including zero touch onboarding of devices
- Provides ultimate flexibility: deployment choice, cloud platform choice, OS choice
- Offers unlimited data duration for more informed networking decisions

Learn more at ExtremeNetworks.com/ExtremeCloud-IQ



 ADVANCE
WITH US

A historic evolution of Wi-Fi technology

Wi-Fi technology is ingrained into our everyday lives. And now, Wi-Fi 6 and Wi-Fi 6E are part of a wireless paradigm shift toward *infinitely distributed* connectivity. Enterprise companies need to connect anybody, anywhere, to any other person, device, or application. Wi-Fi 6 technology focuses on high efficiency and Wi-Fi 6E brings us 1200 MHz of new 6 GHz spectrum to deliver an enhanced *consumer-centric* experience in all enterprise verticals.

Inside...

- Understand current Wi-Fi challenges
- Leverage the 6 GHz spectrum bonanza
- Implement security in a Wi-Fi 6E world
- Provide real-world deployment solutions
- Learn how OFDMA improves efficiency
- Implement downlink and uplink multi-user (MU) communications



ADVANCE WITH US

David Coleman works at the Office of the CTO for Extreme Networks. He is the author of numerous books, blogs, and white papers about Wi-Fi and wireless networking.

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

for
dummies®
A Wiley Brand

ISBN: 978-1-119-80787-2



WILEY END USER LICENSE AGREEMENT

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