White Paper

What to know about Wi-Fi 7

June 2023

Table of Contents

Intended Audience

This document addresses factors and questions related to Wi-Fi 7. As of this writing the standard has not been ratified by the IEEE nor certification defined by the Wi-Fi Alliance so some factors may change between this publication and ratification.

This document is written for and intended for use by technical engineers with some background in Wi-Fi and 802.11/wireless engineering principles.

For more information on how to configure RUCKUS products, please refer to the appropriate RUCKUS user guide available on the RUCKUS support site. https://www.ruckusnetworks.com/support/

Document History

Revision number	Date	Content developer	Reason for revision
1.0	19 June 2023	Jim Palmer	Initial publication

Glossary

- Access Network Query Protocol (ANQP)

- Automatic Frequency Coordination (AFC)

- DFS – Dynamic Frequency Selection; utilized in Wi-Fi 5 GHz operation to vacate a channel if an incumbent signal is present, usually radar.

- EIRP – Effective Isotropic Radiated Power; hypothetical power radiated by an isotropic antenna to achieve the same signal strength. https://en.wikipedia.org/wiki/Effective_radiated_power

- Fast Initial Link Setup (FILS)

- Frame – a Layer 3 packet encapsulated with a Layer 2 header

- IEEE – Institute of Electrical and Electronics Engineers https://www.ieee.org/about

- Jitter – A variance in latency, or the time delay between when a signal is transmitted and when it is received; measured in milliseconds (ms).

- Latency – The delay in network communications; the time that it takes data to transit a network.

- Low Power Indoor AP (LPI)

- Over the air (OTAR)

- Preferred Scanning Channel (PSC)

- QAM – Quadrature Amplitude Modulation

- QoS – Quality of Service; the act of adding priority tags to a frame or packet to increase the likelihood of it being delivered first based on a priority ranking system.

- RF – Radio Frequency, used at Layer 1 for wireless communications.

- RSSI – Received Signal Strength Indicator; estimated measurement of how well the signal is received.

- RTP – Real-time Transport Protocol; a network protocol for delivering audio and video over IP networks.

- SNR – Signal to Noise Ratio; an absolute value of the difference between the signal received and the noise floor at the time. Value is always represented as a positive number.

- Standard Power AP (SP)

- U-NII Band – Unlicensed National Information Infrastructure RF spectrum for Wi-Fi 6E

- U-NII-5 through 8; 5.925 GHz to 7.125 GHz

- Very Low Power AP (VLP)

- WFA – Wi-Fi Alliance https://www.wi-fi.org

Overview

If it doesn't feel like it has been very long since the introduction of 802.11ax and Wi-Fi 6E, it's because it hasn't. When we examine the previous timelines of the PHY amendments to the IEEE 802.11 standard, it's usually about 4 years or so between amendments. However, there are two specific reasons why it feels really fast for 802.11be (Wi-Fi 7) being released.

802.11ac was released in two waves, but not much fanfare was made about the second wave. The only difference with Wave 2 was the introduction of 160 MHz channel bandwidth. The chipsets at the time weren't available at the initial release to support that channel bandwidth, so that feature came much later. Since 160 MHz channel bandwidths are impactable for most designs, the impact of Wave 2 was minimal.

When we look at the PHY amendments that came before 802.11ac, they were being released at a rate of about every 4 years. However, the timetable between 802.11ac and 802.11ax was seven years. In the world of technology, seven years is a LONG TIME. Sure, there was 802.11ac Wave 2, but as mentioned, this was not a huge development that many people were waiting on.

When compared to 802.11ax (Wi-Fi 6) and the release of its "Wave 2" which was the 6 GHz band and Wi-Fi 6E, you can't compare the waves between the two. With seven years of not much to announce between Wi-Fi 5 and Wi-Fi 6, and then four years between Wi-Fi 6 and Wi-Fi 7 with the huge announcement of 6 GHz in the middle, for many it has the feeling that every time they turn around there is a new technology being pushed on them, and it can be overwhelming.

Do we really need Wi-Fi 7?

For many people, this is a legitimate question. With the Fear Of Missing Out (FOMO) being amplified by the many social media avenues present today, this question is being asked more and more. The short answer is yes, we need Wi-Fi 7. Even as we are still learning about 802.11be (the IEEE name for Wi-Fi 7) we are already seeing information about 802.11bn, which is the next PHY rate amendment to the 802.11 standard. The Wi-Fi Alliance is the organization that assigns the "Wi-Fi" name to things so for now we will use the IEEE name instead while assuming that eventually 802.11bn will be named Wi-Fi 8.

It's already on the horizon, and given the timelines looking back in history, it is pretty safe to say that it will be here before we know it, and while we will still be talking about Wi-Fi 8 we will start hearing about Wi-Fi 9. At this point it is just simple counting and progression of technology that pretty much ensures that there will be a Wi-Fi 9, and then a Wi-Fi 10 after that.

This can be debated as much as we want, but it has never slowed the progress of technology in the past so we better embrace that Wi-Fi 7 is coming, and yes, there will be generations beyond that. To better understand why we need Wi-Fi 7, we need to look past the PHY rate which normally defines the generation of Wi-Fi being discussed. The table in the following section helps to understand the relationship between the PHY and the Wi-Fi naming conventions.

Wi-Fi Progression

As technology progresses forward, it helps to compare it with previous generations to fully understand the primary differences in this forward progression and what we gain as the technology advances, and the primary drivers for that gain.

802.11	Suffix	Wi-Fi Alliance Marketing Term	Maximum Speed (PHY)	Year Introduced	Supported Bands
802.11a	N/A	Wi-Fi 2*	54 Mbps	2000	5 GHz
802.11b	N/A	Wi-Fi 1*	11 Mbps	1999	2.4 GHz
802.11g	N/A	Wi-Fi 3*	54 Mbps	2003	2.4 GHz
802.11n	HT	Wi-Fi 4	600 Mbps	2008	2.4 & 5 GHz
802.11ac	VHT	Wi-Fi 5	6.9 Gbps	2012	2.4 & 5 GHz
802.11ax	HE	Wi-Fi 6, Wi-Fi 6E	9.6 Gbps	2019	2.4, 5, & 6 GHz
802.11be	EHT	Wi-Fi 7	46.1 Gbps	2023	2.4, 5, & 6 GHz

Table 1: Wi-Fi Naming Convention

Note: The Wi-Fi Alliance only officially named Wi-Fi 4 (802.11n), Wi-Fi 5 (802.11ac) and Wi-Fi 6 (802.11ax). Wi-Fi 1 – 3 is the assumed naming convention.

In addition to seeing the increase in speeds, it is also important to pay attention to the suffix that has now become the standard for the PHY amendments. By paying attention to this, it gives us an idea of what the IEEE is focusing on with that standard, i.e., why did they make the choices they made for that generation.

These suffixes are as follows:

- HT – High Throughput. Wi-Fi 4 is 802.11n HT

- VHT – Very High Throughput. Wi-Fi 5 is 802.11ac VHT

- HE – High Efficiency. Wi-Fi 6 is 802.11ax HE

- EHT – Extremely High Throughput. Wi-Fi 7 is 802.11be EHT

When viewing the table, it is easy to see why HT and VHT are named that. Going from a maximum rate of 54 Mbps to 600 Mbps is a big leap in throughput speeds. Going from 600 Mbps to 6900 Mbps is a tenfold increase, so Very High Throughput makes sense. When we examine the speeds for 802.11ax HE we do see an increase in speed, but in fact the goal of that amendment can be seen in the suffix name – High Efficiency. When realizing that this was the introduction of OFDMA modulation, it makes more sense. While there was a speed increase, that wasn't the overall focus.

When we look at the suffix for Wi-Fi 7 – Extremely High Throughput (EHT), we do indeed see the speed increase that one would expect, going from 9.6 Gbps to a PHY maximum of 46.1 Gbps. When looking forward to the next

amendment in the works, it is easy to tell from the name that it won't be focused on speed but on stability and reliability.

Using this pattern, it is easy to see that every other generation will focus on speed, and the intervening generation will focus on refining the 802.11 protocol to meet the ever-changing needs of wireless technology. Even though Wi-Fi 7, 802.11be EHT, is focused on speed, there are other enhancements that are being introduced with Wi-Fi 7 that will allow engineers and developers alike to refine the enhancements coming in 802.11bn, which is generally accepted will be named Wi-Fi 8 when it is released in 2028 or 2029.

While networks of today might not need the 46.1 Gbps of throughput forecasted to be possible with Wi-Fi 7 (but very unlikely to be achieved in a real-world scenario; see the section on Extremely High Throughput) there are other features in Wi-Fi 7 that will help networks of today that support Real Time Protocol (RTP) applications like Augmented Reality/Virtual Reality (AR/VR), gaming, and any other low latency applications that rely more and more on wireless connections in order to operate.

Wi-Fi 7 Enhancements

As with every generation of Wi-Fi technology, there are some key features that end up being the focus of that release and Wi-Fi 7 is no different. While there are some other features that might make it through the ratification process (expected to happen near the end of 2023) there are several that vendors are looking at as the primary enhancements with Wi-Fi 7.

• Extremely High Throughput

• Punctured Transmission

• Multi-Link Operation

• Enhanced Quality of Service

Each one of these enhancements can function on their own, but it is expected that a couple of them will work in tandem with the other to offer an even bigger improvement with this generation of Wi-Fi.

In addition to these Wi-Fi 7 specific enhancements, there is a development that underpins most of these new features, and it is in the form of the new 6 GHz spectrum. To ensure a consistent understanding of 6 GHz we are going to review this new spectrum before these new Wi-Fi 7 enhancements.

6 GHz and Wi-Fi 7

Released in 2020 in some regulatory domains, the 6 GHz spectrum was announced with Wi-Fi 6E. One of the bases for the Extremely High Throughput designation for Wi-Fi 7 is the ability to run wider channel widths of 320 MHz which only works with the inclusion of the 6 GHz band.

In addition to Wi-Fi 7, it is expected that future generations of Wi-Fi will rely heavily on this new spectrum so basic understanding of 6 GHz is required. While there is a lot more than goes into the operation of 6 GHz, here are some basics to know.

6 GHz devices are full of new and innovative features. Since there are no other legacy Wi-Fi devices in the 6 GHz spectrum, many of the old communications rules are no longer needed. New security and traffic protocols specific to Wi-Fi 6E help to keep client traffic secure and speed up the AP discovery process, some of which are listed below.

- **WPA3** – Over the air (OTAR) communications are secure even before authentication takes place.
- **Faster AP discovery** – Every 4th 20 MHz channel is a preferred scanning channel (PSC).
- **Reduced probe request** – Clients are not allowed to probe in a PSC channel unless an AP has been discovered.
- **Shortened beacon size** – Since there is no need for backwards compatibility the beacon does not need to accommodate legacy device information, a 6 GHz beacon contains 13 separate sections as compared to a Wi-Fi 6 5GHz beacon contains 17 sections.
- **Multiple SSIDs in a single beacon** – This has been around for a while but was never supported on the client side, but since 6 GHz will be greenfield all of the new clients will support this. This allows APs to broadcast up to 4 SSIDs in one beacon as opposed to the current method of one SSID per beacon.
- **Unsolicited probe responses** – Standard legacy beacon intervals are 102.4 msec, unsolicited probe responses can be transmitted every 20 msec, this allows a client device to decide whether the AP is suitable for a connection through passive scanning rather than active probing as a client can now listen for just 20 msec. There is also no frame exchange during this process, reducing contention.
- **Fast Initial Link Setup (FILS)** – These contain mini beacons that allow a client to determine which AP to connect to, these are commonly referred to as "short SSIDs".
- **Out of band AP discovery** – This feature allows client devices that are connected with 2.4 or 5 GHz to discovery 6 GHz APs without turning on their 6 GHz radio.
- **The Reduced Neighbor Report (RNR)** – Wi-Fi 6E and Wi-Fi 7 clients associated in 5 GHz can request an RNR to identify APs transmitting in 6 GHZ and then move to that AP or radio. (Part of Out of band AP discovery)
- **Access Network Query Protocol (ANQP)** – This capability is used in multiband networks; the APs would send ANQPs separate from probes or beacons, this is more efficient and allows clients to discover 6 GHz APs without leaving their current connection. (Part of Out of band AP discovery)

Equipment Classes

Low Power Indoor (LPI) APs

Low Power Indoor, this will be the most common class of AP, most home and enterprise APs will fall into this classification. LPI APs can operate across the entire band as their power is considered safe for incumbents and the structural materials, they are deployed inside of shield from RF leakage.

These cannot be used outdoors, and the FCC has regulations to prevent just such usage.

- No external antennas.
- No battery power.
- No weatherized enclosures.

The actual power level (EIRP) for LPI APs is not defined in absolute dBm, as for the lower bands, but at 5 dBm/MHz, adding 3 dB for every doubling of channel bandwidth, which gives 18 dBm EIRP for a 20 MHz channel, and up to 27 dBm for a 160 MHz channel. This is possible due to incumbents operating in a very narrow band. It is advantageous to the Wi-Fi network because background noise increases proportionally with bandwidth, so the SNR for a Wi-Fi receiver will be constant for different channel widths, given maximum transmit power levels.

NOTE: The European regulatory authorities allow a max EIRP of 23 dBm.

Standard Power (SP) APs

Standard Power APs are defined as any Wi-Fi 6E AP operating either indoor or outdoor operating at power levels above the LPI AP rules.

These APs are where the Automatic Frequency Coordination, or AFC, comes in. Operating at power levels above that of a LPI AP, SP APs can interfere with incumbents and must be able to check in with the AFC to comply with FCC regulations.

The maximum allowed power for SP APs is set at 36 dBm EIRP. (USA only)

SP APs are only allowed in the UNII-5 and UNII-7 bands which allows for forty-one 20 MHz channels; twenty channels with 40 MHz wide; nine channels at 80 MHz wide, or four channels at 160 MHz.

NOTE: No SP APs are approved for use by the European regulatory authorities and only LPI APs are allowed with a max EIRP of 23 dBm.

Client devices

Client devices are assumed to be connected to a Wi-Fi 6E AP operating within regulations.

Client devices must connect at 6dB lower than the AP they are connecting to.

Very Low Power (VLP) devices

These are defined as hotspots, wearable devices and vehicular mounted devices transmitting at very low power levels. These are expected to cause no interference issues and will not require AFC connectivity and will be able to utilize all the UNII bands allocated for Wi-Fi 6E.

All of the standards for this class have not been defined as of yet.

Wi-Fi 7 Security

Information security, and network security, are always a big topic so to help explain security as it relates to Wi-Fi 7, we are going to take a second to cover this question.

WPA3

WPA3 is a standard that has been defined by the Wi-Fi Alliance and has been with us since 2018. Just as with previous versions of WPA security, just because a newer version has been released doesn't mean that the previous versions have been deprecated. The one exception to this rule is the new 6 GHz spectrum. Many times, WPA3, Wi-Fi 6/6E, and Wi-Fi 7 gets intermixed, but it's more nuanced than that.

The 6 GHz spectrum has a mandate, no matter the IEEE or WFA naming of the devices using that spectrum, that WPA3 or OWE is mandated. This will be either WPA3-Enterprise or WPA3-SAE. Also, in 6 GHz, open networks aren't allowed so when any 6 GHz Wi-Fi signal is seen, it will have some type of encryption in place. To be clear, this is a spectrum mandate, not a Wi-Fi generation mandate.

WPA2

While WPA2 has been superseded by WPA3, it doesn't mean that WPA2 is dead. When searching the internet for information, many charts can be found that show Wi-Fi 6E and Wi-Fi 7 with WPA3 security, but these charts fail to address this nuanced discussion. Wi-Fi 6E is focused just on the 6 GHz spectrum that was released after Wi-Fi 6, but before Wi-Fi 7, so Wi-Fi 6E is mandated to only support WPA3. What is missing in this conversation is that Wi-Fi 6 and Wi-Fi 7, while supporting WPA3, there isn't a mandate that WPA3 is the only security protocol used with these generations of Wi-Fi.

Wi-Fi 7 does in fact have a mandate that it is backwards compatible with previous generations of Wi-Fi, so it only makes sense that while WPA3 will be supported, it will still be possible to deploy a Wi-Fi 7 network and still be able to support legacy clients on legacy bands, using WPA2 as an encryption option. While we would always like to only deploy the latest and greatest technology, there is always a need to support devices in networks that aren't the latest and greatest technology, but for the operation of the business, usually end up being the most critical of devices for business operation.

While network operators should always look for the opportunity to advance their network configurations to remove older technologies (802.11b and TKIP being two old technologies that should be disabled at this point) there will always be this balancing act between the latest technologies and supporting all the devices needed to access the network.

The one beacon of hope that Wi-Fi professionals are looking at now is if a device will now support 6 GHz, by extension it then has to support WPA3, which allows network architects to design networks that will only operate with WPA3. Of course, this thought process relies heavily on two key aspects that we are still waiting on.

1. The WFA will make 6 GHz support mandatory for gaining Wi-Fi 7 certification.

2. All of the devices that will be operating on the network are Wi-Fi 7 certified or supports Wi-Fi 6E.

If and when this happens, we will finally start to see WPA2 security start to be retired in favor of WPA3, which is a monumental leap forward in securing our networks and our data.

Extremely High Throughput

Seeing as this is the primary naming convention for the new Wi-Fi 7 standard, we should cover this one first. There are some straight forward changes that enables the high throughput number in Table 1 above, but all of these will need to be in place to see those numbers. As most engineers will be aware of, it isn't always possible to achieve the perfect scenario, especially in the wild, so your numbers will more than likely vary from those mentioned here.

Increased Quadrature Amplitude Modulation

Quadrature Amplitude Modulation, or QAM, is how Wi-Fi devices encode data to be transmitted over the air. Starting way back at 16 QAM, each generation of Wi-Fi has seen an increase in the QAM capability, with Wi-Fi 7 introducing 4096, or 4K, QAM. Not to be confused with 4K streaming on your TV, this 4K QAM is the ability to represent 4096 data points in what is known as a constellation, but 4096 QAM comes with a downside.

To achieve this high QAM rate, the RF environment needs to be clean. The margin of error when using a QAM rate this high is miniscule and in unlicensed spectrum, the challenge is even more intense. Some numbers that engineers will be looking at to achieve this high modulation rate will be something along the lines of:

- 40 to 45 dB Signal to Noise Ratio (SNR)
- -43 to -40 dB Receive Signal Strength Indicator (RSSI)

In certain deployments this might be achievable, but in high density deployments or in shared RF spectrum, this might be a bridge a little too far to achieve this modulation.

Increased Maximum Spatial Streams

As with 4K QAM, the number of spatial streams that are "available" has increased from 8 in Wi-Fi 6 to 16 in Wi-Fi 7. Also, as with QAM, this also comes with a caveat. While there is the possibility[1] that there might be an AP that is released that features 16 transmitters, 16 receivers, and 16 spatial streams (written as 16x16:16) there isn't any expectation from anyone that we will see anything close to this on the client side. This is important when executives get their brand-new Wi-Fi 7 device and want to run a speed test on it.

While most people understand that Wi-Fi adds in some management overhead that keeps us from burying the speed test needle at 46 Gbps, keep in mind that 46 Gbps number is only applicable with both sides of the link are 16x16:16. Most mobile devices would drain their battery in a matter of minutes if they tried to support that amount of RF technology in a mobile device, so we still expect to see these mobile devices running a 2x2:2 radio configuration at best.

While nothing to sneeze at, it means that for the most part, these spatial streams will be the limiting factor in drag racing this new technology.

Standard	Wi-Fi 6/6E	Wi-Fi 7
Max Speed with 1x1:1	1.2 Gbps	2.9 Gbps
Max Speed with 2x2:2	2.5 Gbps	5.8 Gbps
Max Speed with 16x16:16	9.6 Gbps	46.4 Gbps

Table 2: Expected Throughput Comparison

Increased Channel Bandwidth

As with Wi-Fi 5 Wave 2 that introduced a ridiculously wide channel (160 MHz wide channel), Wi-Fi 7 is also introducing a channel width that makes engineers scoff, and that is 320 MHz wide channels. The only difference

---

[1] **While a possibility, it is highly unlikely there will ever be an AP with 16x16:16.**

now is at least this time around there is spectrum to support this new channel width with the 6 GHz that came with Wi-Fi 6E.

As with most of the technology, this increased channel bandwidth wouldn't have been possible without the introduction of a feature or technology from a previous generation.

Now, granted that not every regulatory domain has 6 GHz yet, and there are still some areas that only have a smaller portion of that spectrum (500 MHz of spectrum compared to 1200 MHz that was released in some areas). However, in areas that gained access to this full spectrum, there is the possibility of running three non-overlapping channels of 320 MHz wide. While only running three channels might not work in most environments, it does leave open some scenarios where one 320 MHz wide channel could be mixed in with the remaining five channels at 160 MHz wide.
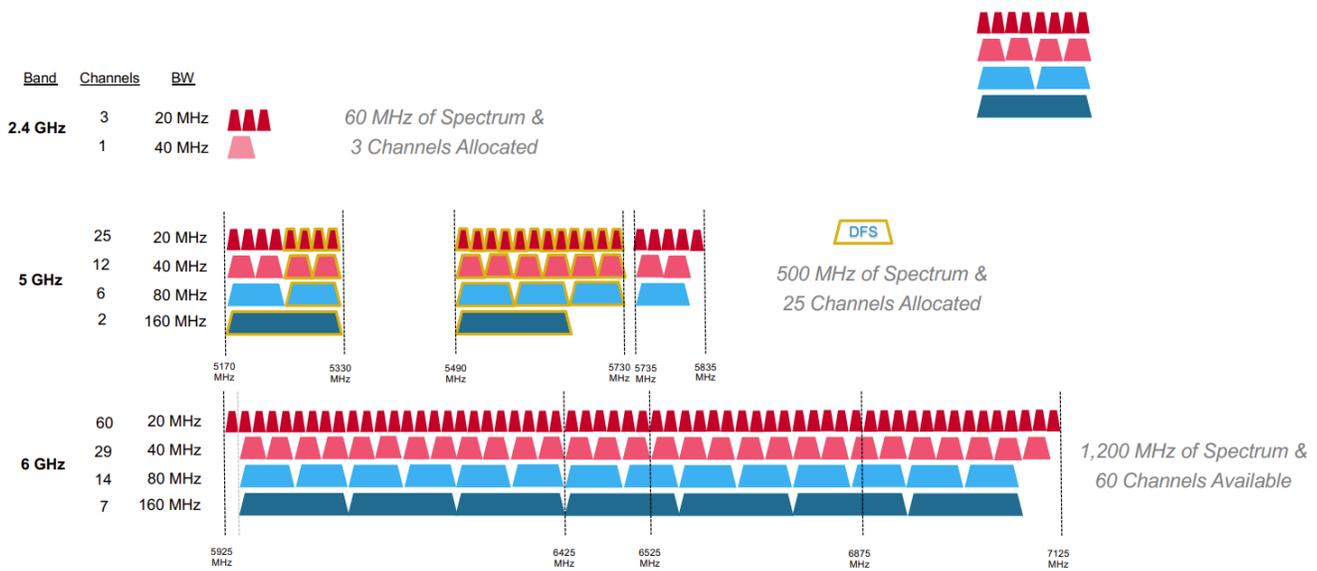

Figure 1: Full Spectrum Allocation courtesy of Broadcom

As with Wi-Fi 6E, 6 GHz is still restricted to Low Power Indoor operation unless connected to an Automated Frequency Coordination system. Only then is standard power operation an option. While there have been advancements to the AFC certification by the United States FCC, as of this publication it still isn't freely available to be used. Current projections are looking at full AFC implementation in the second half of 2023.

Wi-Fi 7 will build upon this work from Wi-Fi 6E and we expect to see much smoother operations in this band as more devices (clients and APs) are released with the Wi-Fi 7 certification from the Wi-Fi Alliance, slated to come near the end of 2023.

Punctured Transmission

Wi-Fi 7 introduces a new terminology that at first glance doesn't sound like it has any place in the Wi-Fi world, but after a little explaining, it is easy to see where this could be a benefit.

The word "punctured" by itself is defined as "to pierce or perforate" or "to make (a hole, perforation, etc.) by piercing or perforating." This terminology just leads to confusion as none of this makes much sense how it can relate to Wi-Fi. The key to understanding what this feature accomplishes is by looking at it not from the aspect of a desired signal, but from the aspect of any unwanted interference on the channel.

All or nothing spectrum utilization

Interference in unlicensed spectrum is something that Wi-Fi professionals have come to accept as just part of doing business. In the past, any interference or unwanted signal that interfered with a desired channel width would result in the stations on the BSS (known as a STA) to simply drop the wider channel and revert back to a smaller channel width to avoid the interference, or in some instances downshifting all the way back to a default channel width of 20 MHz.
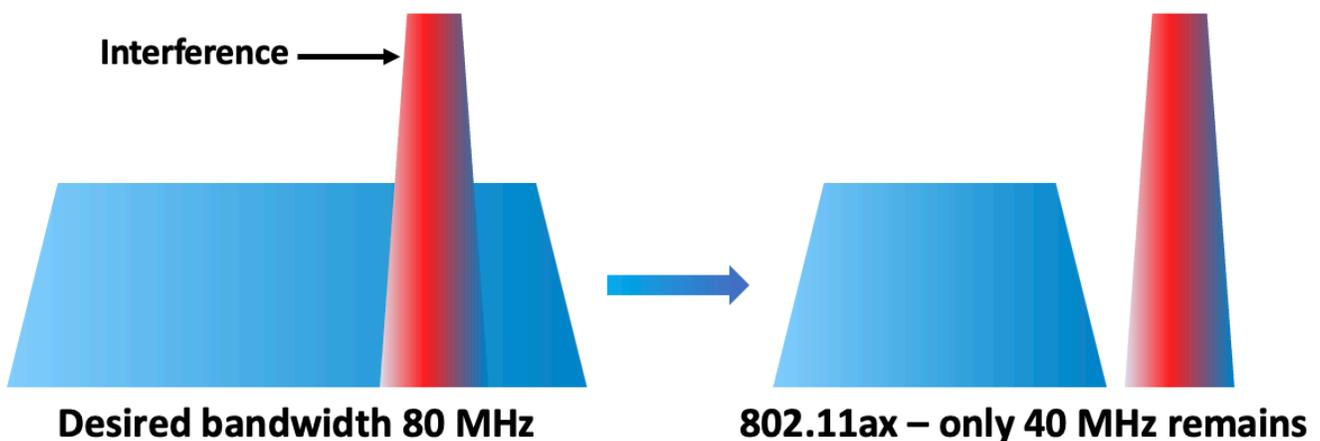

Figure 2: Interference on a Channel

As discussed earlier, wider channels are key to achieving faster speeds. As devices move data faster, they spend less time moving that data, which then frees up the channel for other devices to transmit their required data. While the end result isn't a "faster" connection, from the end users viewpoint the service is indeed faster, as they got their data faster than they would have.

In the past, when interference happens there was no flexibility in how the channel width would downshift. Any interference, even a narrow band interferer, could render the rest of the spectrum unusable due to the rigidness of the rules.

Flexible Spectrum Utilization

Using the image above, imagine the red interfering signal as a puncture in the spectrum. With punctured transmissions, there is now a flexibility introduced into the standard in how devices will deal with that puncture. Instead of simply discarding everything else as a result of this puncture, this enhancement allows the channel to "heal" around this puncture and still be able to operate in the unaffected spectrum.
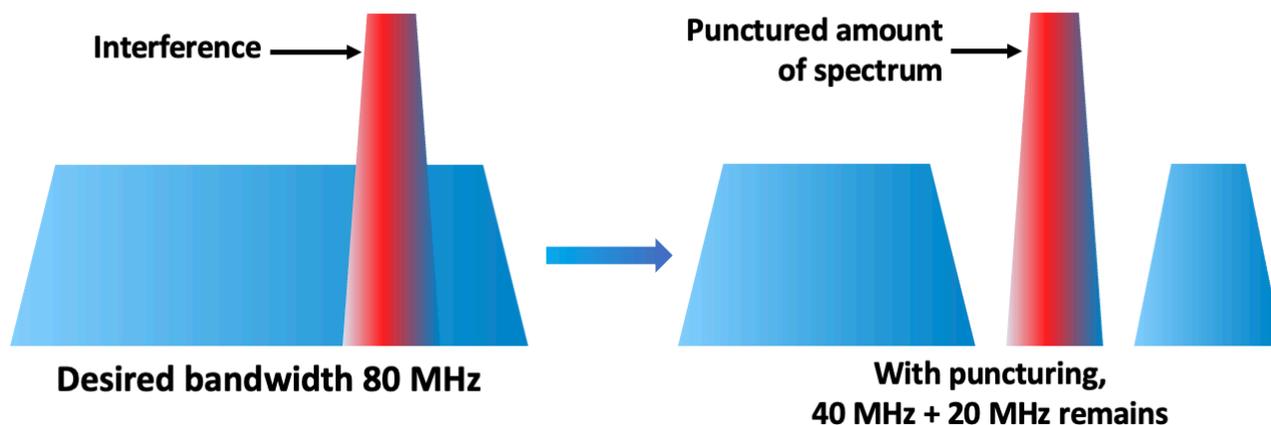
Figure 3: Punctured Transmission on a Channel

While throughput will take a hit as the channel downshifts from 80 MHz worth of bandwidth, it is still able to take advantage of the remaining, unaffected spectrum. In the example above, 20 MHz worth of spectrum is added back to working channel to be utilized, resulting in a decrease of only 25%, not the 50% or even 75% that was seen previously. With this added flexibility in spectrum and channel management, the resulting system has a natural increase in channel availability and conservation of throughput which leads to lower latency on the BSS.

When viewed from the aspect of newer clients and demands on the networks of the future, technologies like Augmented Reality (AR), Virtual Reality (VR), and immersive esports which are very sensitive to latency and jitter, preserving as much spectrum to be used to transmit data only makes the most sense.

Punctured transmission outdoors

Another place where punctured transmissions has the potential to be useful is in outdoor 6 GHz operation. Outdoor 6 GHz operation utilizes a service called Automated Frequency Coordination (AFC) to protect the existing operation of incumbent users of the spectrum.

In previous generations of Wi-Fi there has been a requirement to abandon the entire channel if there is any incumbent operation detected (referring to DFS operations and radar in 5 GHz). With punctured transmissions, the opportunity now exists, theoretically at least, that the incumbent operation could still be protected without the need to abandon the entire channel.

As of this publication this is still a theory as AFC operation is still in testing and final functionality is still to be determined. To learn more about 6 GHz operation you can visit here https://www.commscope.com/solutions/enterprise-networks/wi-fi-6/ .

Multi-Link Operation

The concept of multi-link operation (MLO) is straightforward, and something that Wi-Fi professionals have been looking forward to many years. One of the key tenets of Wi-Fi is it Carrier Sense Multiple Access – Collision Avoidance (CSMA-CA) and as such, operates at only half-duplex. This means that at any given time, on any given

BSS, only one device is transmitting at a time. That means a client device like a tablet or phone is either transmitting or receiving at any given time, but not at the same time.

In an earlier draft of the 802.11be standard there was a concept known as "Coordinated Access Point Transmission" that was going to be part of this feature, but it was moved to the next standard that we expect to be Wi-Fi 8 (802.11bn). Whether it makes it through that process is yet to be seen, but for now MLO in Wi-Fi 7 is still limited to a single device talking to a single AP, and still either transmitting or receiving, but not both. The draft option that was pushed to a later release does offer a glimpse of a time when Wi-Fi devices might make the jump from CSMA-CA and simplex operation to something more. What that will look like is still unknown, but it does give us hope for new developments in the future.

As with the 320 MHz wide channels, MLO is another feature that is building upon the availability of the 6 GHz spectrum. If the options were to only utilize the 2.4 GHz and the 5 GHz spectrum, the usefulness of MLO quickly dwindles. When we consider the ability to use both the 5 GHz AND the 6 GHz band, either in conjunction or for different purposes, the usefulness of MLO can be quickly seen.
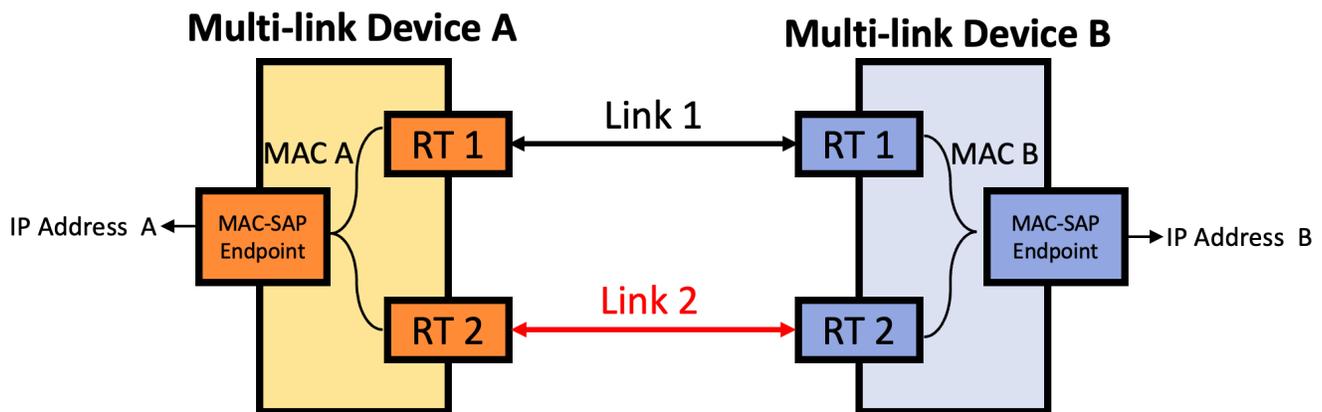

Figure 4: Multi-Link Operation

With new Wi-Fi 7 devices being able to support all three bands, and the possibility of some instances needing flexibility to pick and choose the band to operate in at that given time, why not use these radios in our devices? The concept here is that each radio in the client device will establish a radio link with the AP, but in different bands. Link 1 could be on a 5 GHz channel while Link 2 could be established on a 6 GHz channel. Each device then coordinating the data that is sent across those links. Think of it as channel bonding as it was introduced in Wi-Fi 4 (802.11n) but across spectrum bands without actually bonding those channels together.

With these links established, there are a couple of options that can happen. Which option is selected, and how often, is still yet to be seen as any Wi-Fi professional will tell you seems to happen when things are deployed into the wild (things like Beamforming and OFDMA comes to mind).

Potential options as how this will be utilized are:

- Link Redundancy
- Link Aggregation
- Link Selection

To be clear, it might be possible to switch back and forth across the links to accomplish one option at a time, but there won't be any combining of these options.

Link Redundancy

Link redundancy is all about making the transmission of the data being sent more resilient. In this option, the data will be duplicated and sent across both links independently. The receiving end then compares the data received and can do an error correction of sorts so that even if one link is degraded, the device can still send the required ACK to acknowledge receiving the data. More ACK's means less retransmissions, which means a more efficient BSS, which means more clients served in the same amount of time, making the BSS feel "faster" even though nothing here changed.

In applications like Real Time Video and Voice (like Virtual Reality) having a resilient link that gets that data through the first time, without any jitter or latency, is critical to that application. The same can also be said for applications like automated manufacturing or remote medical procedures. When it absolutely has to make it there, link redundancy with its added resiliency is the answer.

Link Aggregation

Link aggregation is as close to the channel bonding example as we can get. In this scenario, both devices in the BSS realize that both the 5 GHz and 6 GHz link are available to use so why not take advantage of it? If a device can achieve a throughput of 2.5 Gbps on a 2x2:2 link on Wi-Fi 6 operating at 5 GHz, and a throughput of 5.8 Gbps on a 2x2:2 link on Wi-Fi 7 at 6 GHz, why not combine those numbers together to achieve a throughput of 8.3 Gbps?

To be clear, this sounds very enticing, but the process to make this happens means there isn't any redundancy on the link. The device will take the total payload to be transmitted and split it up. Some will be destined for the 5 GHz link and the rest will be sent over the 6 GHz link. If any frame doesn't make it across the link, then that frame will have to be retransmitted, adding to the retry percentage on the BSS and in turn slowing down the other clients while this data is retransmitted.

While the improved throughput could offset a small number of retransmissions, allowing for a higher overall performance rate of the BSS, these types of questions are still waiting to be answered. Can the devices keep up with this? How will it make the decision to use link aggregation? At what point will devices stop using link aggregation and go back to a different approach?

We won't be certain of these answers until there are more real-world examples of Wi-Fi 7 networks in operation in the wild. Don't ignore this feature, but at the same time, look at it with a grain of salt.

Link Selection

Link selection, of the three possibilities of MLO, is the more likely option we will see the most of in the wild. In link selection, the stations evaluate each radio link and then decide which link has the best possibility of getting the data across it. By actively switching between which band the data is sent on, the device reduces the utilization on the unselected link by not sending any data across it while sending the data across the link that has the best chance of a successful delivery.

This option best represents the methodologies of the past that have been tried and tested, while still giving the devices the flexibility needed in unlicensed spectrum that can change without any warning. When we start to see Wi-Fi 7 networks in the wild and look at MLO, expect to see this being the most common implementation of the possibilities. Link selection also points to the fact that our Wi-Fi clients will finally start to pay attention to the quality of the link and not just simple items like channel bandwidth or the RSSI value.

When looking forward to Wi-Fi 8 we finally see hope that client devices are finally starting to see the value of listening to the information that has always been present from the AP and working with the Wi-Fi network infrastructure instead of against it. If you have ever worked on troubleshooting Wi-Fi devices in the past, you know the pain of Wi-Fi clients having the final say in all decisions.

MLO Security

Mixed into this enhancement is a development that will improve MLO operation. Previous generations of Wi-Fi utilized a single MAC address for each radio and as such, when creating the encryption keys for a link, needed to create new keys for each link using a four-way handshake. With MLO, there is an introduction of a "higher level" MAC Address that is used across the three radios in a Wi-Fi 7 device.

This single, high level MAC address is used for the encryption keys meaning that as a device switches bands during MLO link selection, a new encryption key doesn't need to be created. With the lower latency requirements of VR/AR/esports, any microsecond saved on the RF side of the transmission is critical to achieve the latency numbers these applications are aiming for.

MLO Summary

At the end of the day, the goal of any Wi-Fi network is to only have to transmit one frame one time. This is the most efficient the network can operate, and the best utilization of RF spectrum. Any network running at peak efficiency will naturally feel faster and better to the end user, and with applications that are now very sensitive to jitter and latency, getting this data through the first time, every time, will be key to successful networks of the future.

Enhanced Quality of Service

The final feature to look at with Wi-Fi 7 is the enhanced Quality of Service (QoS). As discussed previously, applications like AR/VR/esports are sensitive to jitter and latency, and this demand is only going to grow, not shrink. While many of these frames aren't large in size, the timeliness of their arrival is critical to ensure successful implementation of these emerging technologies. As developers push the envelope of what is used for, there needed to be a way to ensure that these critical frames are delivered without any delay.

QoS is good but not great

While there has always been Quality of Service (QoS) tagging in networking, it has always been a reactive approach to the problem with no ability to be proactive. With different levels of QoS, the timers that are required to transmit a frame is reduced, making it more likely that frame will be transmitted sooner than a frame with a lower priority tag.

Where this becomes a problem is when every frame is prioritized as a high priority frame, even when it isn't. The other issue is that even with the highest priority QoS tag applied, there might already be lower priority data being transmitted when the higher priority of data is ready. Wi-Fi works using a collision avoidance algorithm so if there is already a resource using that medium then everything else must wait for that to be completed before anything else can be transmitted.

Enhanced Quality of Service

Realizing that these new technologies need a better way to optimize traffic to decrease latency and jitter, Wi-Fi 7 has introduced "Enhanced Quality of Service" features that should help reduce complaints when it comes to these Real-time Transport Protocol (RTP) applications.

Using MLO for QoS

One feature of this enhancement has already been discussed, and that is MLO using link selection. By being able to examine each link and select the best band to be used (2.4, 5, or 6 GHz) the transmitting device has already taken the first step in ensuring the quality of service for the information being sent.

In addition to feature, there is another feature that was previously introduced but underutilized in Wi-Fi and that is Target Wait Time, or TWT. In Wi-Fi 7, this has been further refined as Restricted TWT.

TWT for enhanced QoS

Target Wake Time was a feature that was originally introduced to the Wi-Fi world in 802.11ah, or unlicensed operation in the 900 MHz band. TWT allows clients to set a target wake time, essentially telling the infrastructure that it was going to go to sleep to conserve battery life of the device, and then tell it when it planned on waking up. This feature then instructed the AP to hold, or buffer, any data for that device until it woke up. To read up more on TWT you can read this blog https://www.commscope.com/blog/2018/802.11ax-fundamentals-target-wake-time-twt/ .

With Restricted TWT, client devices are now able to build upon the original TWT operation but instead of merely informing the infrastructure when it plans to wake up to receive any buffered data, client devices can reserve the channel for future operations, thereby restricting the channel during that time for it's transmit opportunity.
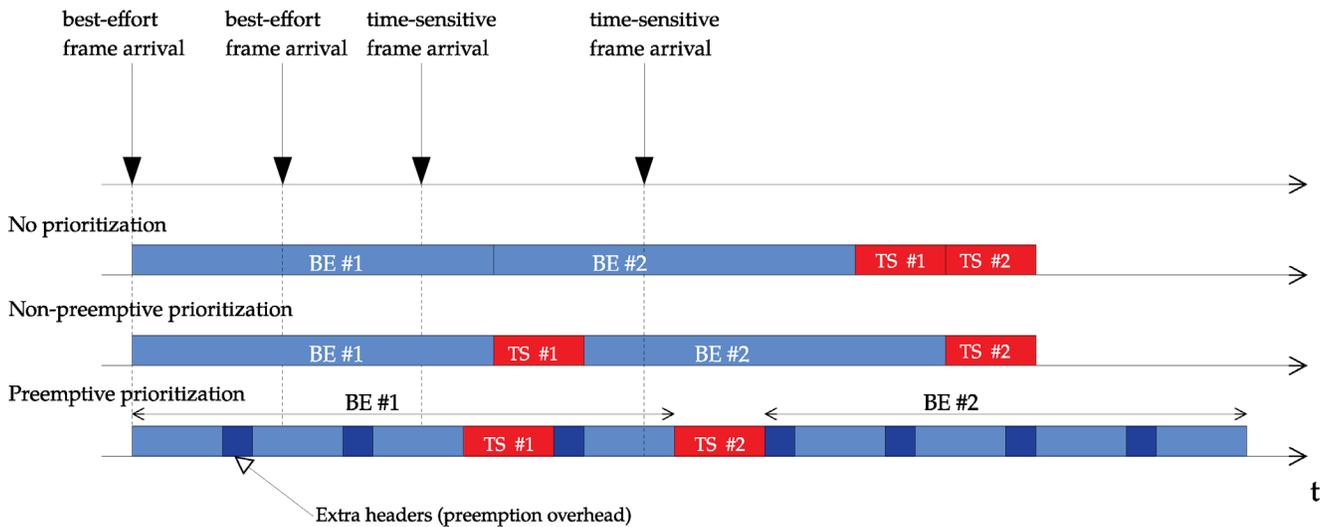
Figure 5: Restricted TWT with Preemptive Scheduling

By adding in extra headers into frames to reserve these future time slots, devices that are supporting RTP applications like voice or video calls, or even medical devices that need to ensure timely delivery of their live-saving data, can reserve a time slot in the future of when they are going to need to transmit their data. This allows the BSS to be very predicable for this type of operation and allow the data to be transmitted exactly when it needs to be transmitted, and not just when the opportunity is allowed.

As with many enhancements, we need to see what is going to be mandated by the Wi-Fi Alliance and what will be optional. If other devices can ignore this feature and send their own traffic instead, then this feature might be relegated to the back burner of the IEEE world. Along with some of the other features we have discussed, this one will be a "wait and see" before we have a celebration for it.

Summary

Wi-Fi 7 promises to be a great step forward in the evolution of Wi-Fi. As our desire for more data to be delivered to more devices, and delivered faster and more reliable than ever before, our ever-increasing mobile world is demanding more from the 802.11 standard. Not to be lost in any of this is the chatter that has already started about 802.11bn (expected to be named Wi-Fi 8 by the WFA).

Just as all previous generations of Wi-Fi were simply steps in an evolutionary chain that has gotten us to 2023, Wi-Fi 7 is still just another evolutionary step in that chain. There will be newer generations that will always be on the horizon promising faster, better, more reliable performance than what we have today. Better roaming, more responsive networks promising an effortless and snappy feel while still offering the flexibility in use that we have today. These evolutions can be a challenge to keep up with for the most seasoned professionals, let alone network administrators and operators who still need to focus on keeping the network running and existing clients functioning.

RUCKUS Networks has always been an industry leader in Wi-Fi and as these generations of Wi-Fi continue to evolve and grow, RUCKUS will always be at the forefront of this evolution and a trusted partner for organizations the world over. To keep up to date with all the different Wi-Fi developments coming now and into the future, rely on RUCKUS to be there at https://www.ruckusnetworks.com .

RUCKUS solutions are part of CommScope's comprehensive portfolio for Enterprise environments (indoor and outdoor).

We encourage you to visit ruckusnetwors.com and commscope.com to learn more about:

- RUCKUS Wi-Fi Access Points
- RUCKUS ICX switches
- SYSTIMAX and NETCONNECT: Structured cabling solutions (copper and fiber)
- imVision: Automated Infrastructure Management
- Era and OneCell in-building cellular solutions
- Our extensive experience about supporting PoE and IoT

RUCKUS®
COMMSCOPE