
TECHNICAL PAPER



With Wi-Fi 6E, Is It Time to Consider a Layered Network Approach?

Chuck Lukaszewski
Vice President & Wireless CTO




TABLE OF CONTENTS

EXECUTIVE SUMMARY	3
ASSIGNING SSIDS TO FREQUENCY BANDS WITH WI-FI 6E	3
THE TRADITIONAL APPROACH	4
EARLY ATTEMPTS AT PARTITIONING 5 GHZ	4
RECENT EFFORTS TO IMPLEMENT LAYERING STRATEGIES IN HEALTHCARE	5
TWO-LAYER PARTITION STRATEGIES FOR 6 GHZ	5
PARTITIONING 6 GHZ INTO THREE LAYERS	7
WHAT OTHER ORGANIZATIONS COULD BENEFIT FROM BAND PARTITIONING?	7
OUTDOOR 6 GHZ	8
HOW DO CLIENTS DISCOVER AND ATTACH TO 6 GHZ ONLY PARTITIONS?	9
BAND PARTITIONING IN COUNTRIES OPENING 500 MHZ ONLY	10
WHAT ABOUT 2.4 GHZ?	10
BEYOND WI-FI: THE MULTI-RAN ENTERPRISE	10
WHAT'S NEXT?	10



EXECUTIVE SUMMARY

A new frontier in wireless is opening, and Aruba will be leading the way with new access points that support the 6 GHz band and Wi-Fi 6E. In the Americas and other countries such as Korea and Saudi Arabia, the 6 GHz band offers 1200 MHz of new, greenfield unlicensed spectrum indoors – a full 3X increase in the amount of Wi-Fi capacity available to enterprises.¹ This creates a fascinating opportunity for wireless network architects who have long wanted to deploy multiple, overlapping Wi-Fi networks to serve different user communities or device classes that have distinct RF performance characteristics. The arrival of Wi-Fi 6E and the subsequent adoption of Wi-Fi 6E devices will alleviate the problem of scarce and congested 5 GHz channels and allow more spectrum abundance. Given the new channels, is it best to continue current practices and deploy all channels in a single layer of coverage? Or, for some enterprises, is it time to start adding at least one additional layer of APs, and segregating the new channels into sub-bands to serve different device types? In such an approach, what are the optimal channel bandwidths for each layer? And what should be the role of the 5 GHz and 2.4 GHz bands going forward in managed networks? This paper explores a potential radical rethinking of how large managed WLANs can be dimensioned and deployed using the 6 GHz band once Wi-Fi 6E capable devices have gained critical mass. It primarily references examples using the full 1200 MHz of spectrum the 6 GHz band, although the concepts can be applied to the 500 MHz markets as well.

ASSIGNING SSIDs TO FREQUENCY BANDS WITH WI-FI 6E

To start, consider how the traditional dual-band enterprise SSID layout could evolve in a tri-band 6E deployment. Table 1 shows the most deployed approach today, while table 2 illustrates a likely enterprise scenario with 6E.

Obviously, Table 1 will vary between organizations. If you don't need to support PSK-based devices, there's no need for that SSID. Or perhaps your PSK devices are specific to one band or the other. Maybe you don't support guest access or do so only with a PSK. And of course, there could be more than one 802.1X SSID, such as in higher education with Eduroam.

Table 2 is where it starts to get interesting. Using the 6 GHz band presents an opportunity for enterprise wireless architects to take a fresh look at their current wireless needs, rather than simply cutting and pasting their 5 GHz configurations into the new band. Given all that new capacity, consider these questions:

- Should you limit 6 GHz to devices belonging to the enterprise or at least those BYOD devices that have been formally onboarded via a Network Access Control (NAC) solution like ClearPass Onboard?
- As the organization invests in native 6E devices for specific business applications, will there be a desire to limit their operation exclusively to 6 GHz to ensure consistent performance?
- Is putting guest traffic in 6 GHz really a good idea – at least until the organization has a better idea of the answers to the first two questions – when 5 GHz provides the widest compatibility for unknown equipment types?
- With three bands available, is it finally feasible to clear 2.4 GHz for exclusive use by all your IoT PSK devices?

With the proposed approach in Table 2, typical CorpNet devices like tri-band capable laptops could be given their choice of 5- or 6 GHz depending on local conditions and driver preferences. Each band would have a single band-specific SSID, including a 6 GHz-only Corp SSID and guest access confined to 5 GHz. Many other permutations are possible. The following section explores options for the optimal channel arrangements in the 5- and 6 GHz bands.

2.4 GHz Radio	5 GHz Radio
Corp_SSID (802.1X)	Corp_SSID (802.1X)
Guest_SSID (Open)	Guest_SSID (Open)
IOT_SSID (PSK)	IOT_SSID (PSK)

Table 1: Conventional Dual-Band SSIDs

2.4 GHz Radio	5 GHz Radio	6 GHz Radio
	Corp_SSID (802.1X)	Corp_SSID (802.1X)
IOT_SSID (PSK)	Guest_SSID (OWE)	Corp_6Only_SSID (802.1X)

Table 2: Potential Tri-Band 6E SSID Layout

¹ Europe has opened 500 MHz in the lower portion of the 6 GHz band. This represents about a 2X increase in capacity in Europe and countries that choose to follow the European approach. This paper focuses on the 1200 MHz Americas approach to show the limit of what is possible.



THE TRADITIONAL APPROACH.




Figure 1: Traditional “Single-Layer” Plan with Each Band Having a Single Pool of Channels

Figure 1 shows the simplest arrangement of the 5 GHz and 6 GHz bands for countries that permit low power indoor (LPI) access to the entire 1200 MHz.

Figure 1 is a “single layer” channel plan, where all the radios offer the same SSIDs. This model simply scales up what wireless architects have been doing for 20 years to include 6 GHz, which is a simple “rip-and-replace” of dual-band APs to tri-band APs featuring 2.4 GHz, 5 GHz, and 6 GHz bands on separate radios. The main difference of course is the ability to use 80 MHz or even 160 MHz as default channel bandwidth in the new band. Note that Figure 1 also includes the new channels of the U-NII-4 band, which has just opened in the US and are supported by Aruba Wi-Fi 6E platforms.² You can clearly see the 3X capacity increase. On the top is the 6 GHz band. This has 14 channels, each of which is 80 MHz wide and permits 1.2 Gbps of peak speed. This “layer” also has a single 40 MHz channel at the very top of the band that is part of the capacity pool. In the middle is the 5 GHz band where 40 MHz channels are in use, which allows for a total of 14 possible channels in 5 GHz with a peak data rate for a 2x2 client of 573 Mbps on each channel. There is even one unused 20 MHz channel shown in grey at the very top. This model will work well for end user types such as general enterprise, K-12 schools, and retailers that want to add a super-fast express lane everywhere and don’t mind co-mingling all user devices.

Let’s consider how the SSID mapping from Table 2 would function under this band plan. The Guest SSID – which is shown upgraded to Enhanced Open (OWE) in transition mode because all 6E APs are required to support this capability – would have guaranteed access to 40 MHz channels. Similarly, the Corp_6Only SSID would have guaranteed 80 MHz channels. And multi-band devices configured for the dual-band Corp_SSID might experience switching between 80 MHz and 40 MHz rates on different bands depending on driver decision making. Perhaps that’s acceptable in most organizations, but what if it isn’t? What if the wireless architect wants to allocate consistent, predictable bandwidths and control which devices use which bands? That brings us to the fascinating new possibility of creating “layers” or sub-bands within either or both large bands.

EARLY ATTEMPTS AT PARTITIONING 5 GHZ

Consider a hospital with multiple distinct classes of devices with widely diverging bandwidth, latency, and life-safety criticality requirements. There have been several attempts over the years to partition the 5 GHz band into two or three sub-bands (or “layers”) to meet these different needs. One well known example was the InnerWireless “HD3” system.

Sub-Band / Layer	Frequency Range, MHz	20 MHz 802.11 Channels
Low	5150 to 5350 MHz	36,40,44,48,52,56,60,64
Mid	5490 to 5590 MHz	100,104,108,112,116
High	5650 - 5850	132,136,140,149,153,157,161,165

Table 3: InnerWireless “HD3” 5GHz Sub-band Channel Assignments

² U-NII-4 spans 5850 – 5895 MHz, allowing for two new 40 MHz or one new 80 MHz channel at the very top of the 5 GHz band. U-NII-4 does not require DFS. Many other countries are evaluating opening this range.



Instead of the usual single AP on a ceiling every ~2,500 ft² or so (250 m²), an above-ceiling enclosure was installed at that same density with three APs in each box. In 5 GHz, this system typically allocated one layer for clinical telemetry data, one layer for general purpose data devices, and one layer for voice. One recommended HD3 channel partition scheme that is presently deployed at some large hospitals in the U.S. is shown in Table 3.

This approach is certainly innovative, but in practice, it was less effective than intended for a variety of reasons:

- Half of the channels in the low and high layers require DFS while half do not, resulting in highly inconsistent roaming behavior within each layer.
- In the middle layer, having only 5 channels is not enough to permit adequate RF spatial reuse, leading to high levels of co-channel interference (CCI).
- All three layers include DFS channels, meaning that fast-roaming voice deployments are challenged. (This can be mitigated with a different partitioning scheme.)
- Guest traffic was limited to 2.4 GHz in the preferred implementation.
- Complex and expensive RF combiners and RF filters were required for each enclosure, along with a complex network of low-loss RF interconnect cables.
- Placing three APs inside a single enclosure with less than 50 cm of spacing between them resulted in high levels of mutual wideband adjacent channel interference (ACI) that filters could not mitigate.

However, this spectrum management strategy was also ahead of its time and had some serious advantages as compared with simply pooling all the 5 GHz channels in a single block as shown in Figure 1. For instance, safety-of-life telemetry could be partitioned from general data and patient (guest) traffic. And voice-over-Wi-Fi, which is widely deployed in healthcare, could be put on dedicated channels without competing with low-rate data. Guest traffic was completely segregated off to the legacy 2.4 GHz band, which left 5 GHz exclusively for the hospital. These advantages rightfully received attention from many hospitals, even if they ultimately chose a more traditional below-the-ceiling deployment.³

RECENT EFFORTS TO IMPLEMENT LAYERING STRATEGIES IN HEALTHCARE

In late 2018, the Center for Medical Interoperability (C4MI) published a technical specification for hospitals called “Trusted Wireless Health.”⁴ C4MI is a non-profit cooperative research and development lab founded by health systems to simplify and advance data sharing among medical technologies and systems. Section 5.1 of the specification deals with AP deployments, and states:

*“APs SHALL be configured such that Clinical Network traffic is not impeded by competing traffic. **This can be accomplished by carrying guest traffic and enterprise traffic on different radios and on different channels...** To avoid interference from overlapping channels, each layer of wireless traffic SHOULD use six or fewer unique channels with a minimum separation of 20MHz between center frequencies.”*

The specification was the consensus product of C4MI’s technical staff plus 16 different C4MI members including such medical device manufacturers GE Healthcare, Philips Medical, Welch Allyn, and Draeger Medical along with electronic medical record providers like Cerner.

TWO-LAYER PARTITION STRATEGIES FOR 6 GHZ

Enter Wi-Fi 6E. How could this design philosophy be applied to the abundant greenfield spectrum in the 6 GHz band for organizations with these requirements?

Figure 2A shows one possible approach where the architect partitions 6 GHz into two discrete blocks of new capacity: Band 6A and Band 6B. Each sub-band uses a different channel bandwidth. Realizing this architecture requires the organization to deploy a new and separate layer of APs and ethernet cabling, but there are so many channels available in 6 GHz that this can be easily justified given the QoS requirements. In contrast, 5 GHz is retained as a single contiguous block of capacity called Band 5. Because Band 5 is not broken out into two blocks, all APs can share a common pool of channels resulting in a doubling of AP density in 5 GHz compared to 6 GHz – which is fine given the large number of channels – and would raise the typical edge data rate in each cell and increase the total aggregate capacity of the layer after accounting for spatial reuse due to wall attenuation.

³ Hospitals are far from the only mission-critical WLAN user that has experimented with these ideas in the two decades since Wi-Fi was first certified.

⁴ <https://medicalinteroperability.org/specifications/cmi-org-twh/CMI-ORG-TWH-D02-20180914.pdf>

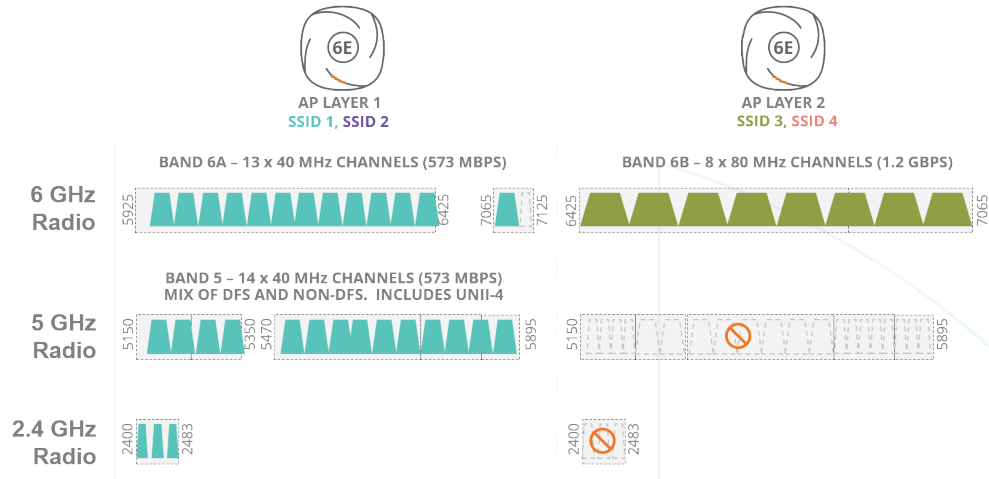


Figure 2A: Band 5 with a Common AP Layer and Band 6 with Two Layers of APs

Let’s take a moment to really reflect on the possibilities just described. Chances are that the hospital network is already running 40 MHz channels in 5 GHz and is heavily congested. With this proposal, the hospital would get one brand new greenfield network with similar capacity, **plus a second greenfield network with double the capacity/bitrate** for super-high bandwidth applications like medical imaging. This strategy allows the hospital to migrate wireless systems on its own schedule to 6 GHz, reserving it for intended users and devices only. And by limiting patient/guest traffic to 5 GHz, it is literally walled off from the clinical IT network capacity, ensuring patients watching Netflix don’t interfere with clinical systems. Because 5 GHz has the widest compatibility, it makes sense that this band remains “open to the public” while the 6 GHz band is reserved exclusively for hospital systems.

Once you start down this rabbit hole, the world starts to look very different, and even more possibilities open up. Figure 2B shows yet another approach to the dual layer strategy.

As long as we are partitioning one band and have deployed two layers of APs at a common density, why not partition the other band? One could combine the 12 x 20-MHz non-DFS channels of U-NII-1, U-NII-3, and U-NII-4 into one block (call it “5A”) for clinical voice/IOT and set aside the 8 x 40-MHz DFS channels of U-NII-2 (“5B”) for patient/guest/public access. Figure 2B shows this layout; the narrow blue channels of Band 5A are non-DFS 20 MHz suitable for voice systems. This layer could be cleared of all other traffic, immediately improving the voice experience through a combination of more channels and cleaner air.

And why stick with just 40s and 80s in the 6 GHz band? Figure 2B shows 20 MHz channels in use for 6 GHz layer 1. There could be a total of 19 such channels – nearly as much as in the entire 5 GHz band. But on top of that, layer 2 is deployed as five 160 MHz channels for massive multi-gigabit data rates, up to 2.4 Gbps for a single station. This could be adjusted per desired outcome, perhaps increasing the number of 160 MHz channels and pulling back the number of 20s in layer 1.

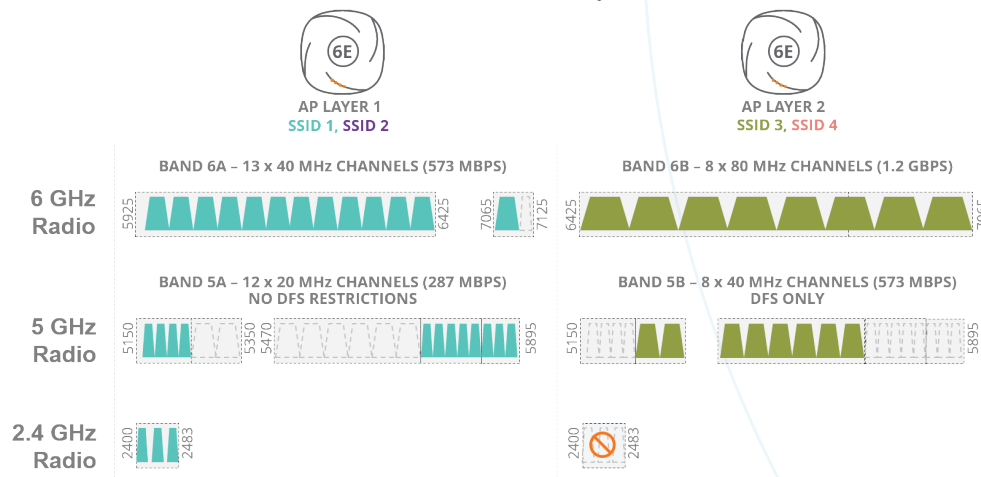


Figure 2B: 5 GHz and 6 GHz Bands Each Partitioned into Two Layers



PARTITIONING 6 GHZ INTO THREE LAYERS

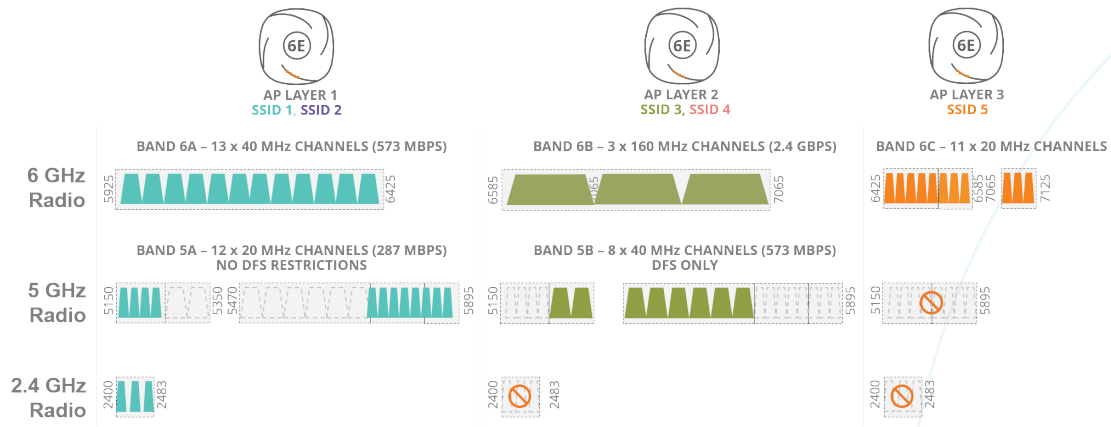


Figure 3: 6 GHz Partitioned into Three Layers with Different Bandwidths (Common 5 GHz Layer)

One reason for adopting spectrum partitioning in 5 GHz is that several important clinical IT systems do not require wide channels. For example, VoWi-Fi handsets work great in 20 MHz, as do many medical telemetry systems. What both have in common is they are dominated by large numbers of very small data packets, with extremely tight latency and jitter requirements. On the other hand, the latest medical imaging systems need massive fiber-like speeds.

Why stop at two layers? For an enterprise that needs to have three completely isolated groups of devices, each requiring different bandwidths, consider Figure 3.

The 6 GHz band is divided into three separate blocks – 6A, 6B, and 6C. 5 GHz is treated the same way as in Figure 3A. The frequency plan splits up the channels like so:

- Band 6A (Clinical Data Network): 6 x 80 MHz, or 12 x 40 MHz channels
- Band 6B (Clinical Imaging): 3 x 160 MHz channels
- Band 6C (Clinical Voice/IOT): 11 x 20 MHz channels
- Band 5A (Clinical Voice/IOT): 12 x 20 MHz channels
- Band 5B (Guest/Public): 8 x 40 MHz channels

Because clinical imaging applications tend to be confined to specific parts of a hospital, one could even envision having two different channel configurations for Band 6B – one where imaging is occurring, and another in a separate area with a more traditional 80- or 40-MHz channel plan.

Achieving this radio design will require three discrete layers of APs and ethernet cabling. In this case, each AP could be in different layers on each band depending on how AirMatch⁵ allocates the channels.

WHAT OTHER ORGANIZATIONS COULD BENEFIT FROM BAND PARTITIONING?

Wi-Fi APs are comparatively inexpensive, both on a per-user and total network basis. Today, the cost of a wall-to-wall WLAN refresh for a campus network is typically 2-3 times less costly than a network-wide edge switching refresh. Due to ethernet port rightsizing over the last decade, there are ample PoE switchports available to support additional APs in most networks. Likewise, cabling costs, while not negligible, are typically less than the AP. In addition on a per-run basis, there has been a long-term trend of large and small enterprises deploying APs more densely with each new generation of Wi-Fi. So as a threshold matter, given that 6 GHz brings such a radical change in capacity, it is entirely appropriate that network architects evaluate the optimal frequency management strategy for their organization. If that leads to a recommendation to double-up some or all of the AP footprint, this can be justified as part of a Wi-Fi 6E migration.

General Enterprise

Some customers have increased their AP densities to as high as 1 per every 1,000 ft² (100m²). There can be many reasons for this, but for general enterprise use, this may well be over-dense. If your network already has above-average density, you may be a natural candidate to consider partitioning with little or no additional AP density increase

Education

Higher education often deploys many APs per room in a lecture hall type setting or in common areas with high densities of users. Educators have long desired to segregate student traffic from other types. Depending on historical

⁵ Aruba's AirMatch capability uses machine learning to automate RF channel assignment, channel width, and radio power assignments.



usage trends, it may be possible to take one AP in each lecture hall and assign it to a different layer for institutional use only without any new cabling or change to AP density. Similarly, in dormitories and student housing, there may be legacy cabling in hallways that predates moving APs directly into student rooms. Such cabling could be repurposed to re-install APs in hallways for institutional users like resident assistants (RAs), security guards, and maintenance teams. In the student rooms, one could even envision a dedicated gaming SSID on its own channel set to provide a more deterministic latency experience for gamers, while at the same time separating this traffic from normal student network access.

Manufacturing

It is not necessary to deploy a second or third layer throughout an entire facility. As noted earlier, manufacturing lines increasingly support very large numbers of different device types. A suitable number of layers could be deployed just on the lines, or even a single line, depending on business needs. For example, it’s increasingly common that Wi-Fi-enabled devices – from simple consumer electronics to appliances and even entire automobiles – pull firmware over Wi-Fi during later stages of the assembly process. Such a manufacturer could deploy one layer of APs for firmware installation, another layer for smart tools, another layer for employee data, and perhaps another layer for robotic AGVs. All with no overlap in their assigned frequencies. With the sheer amount of spectrum in 6 GHz, network architects now have options that never existed before.

Large Public Venues

Stadiums, arenas, and airports are also great candidates for partial segmentation. Due to spectrum limits in 5 GHz, the best practice has been to converge “back-of-house” applications like ticketing, concession point-of-sale, mobile security PDA, press, team data tablets, and even visiting concert acts onto the same very limited set of channels.

With 6 GHz, it becomes possible to subdivide those users into groups and position dedicated new APs where they will do the greatest good. A layer of press APs need only be deployed in the press box, briefing rooms, and on sidelines. Likewise, ticketing and POS are in specific locations. A question that will take a few years to answer is “what is the right channel bandwidth for fans in 6 GHz”? Regardless, it’s unlikely the entire 1200 MHz is necessary to allocate to fans.

Specialized Scenarios

Very high security environments are another natural candidate for segmentation. High-tech, pharmaceutical, or financial services firms all operate labs or trading floors where there is a longstanding desire to physically separate users and not permit certain networks outside of certain areas. Likewise, military and intelligence community users that operate SCIFs or other sensitive areas may desire to have dedicated spectrum allocated for those uses that is used nowhere else.

Some organizations may experiment with using indoor mesh as a cable alternative, recognizing that low power indoor (LPI) rules prohibit the use of external antennas in 6 GHz. But one could envision splitting the band into a dedicated mesh sub-band and a second service sub-band for general usage.

OUTDOOR 6 GHZ

This entire paper has focused on low power indoor (LPI) APs. As of this writing, only two countries have made 6 GHz available outdoors – the U.S. and Canada – and those countries will require a spectrum database solution called an automatic frequency coordinator (AFC). For this reason, Aruba tri-band APs will initially support only LPI in 6 GHz, as it will take a few more years to complete the technical and regulatory work to permit AFC operations to commence.

That said, it is virtually certain that outdoor 6 GHz networks, when they arrive, will use different channel plans than indoor 6 GHz networks. There are two reasons for this.

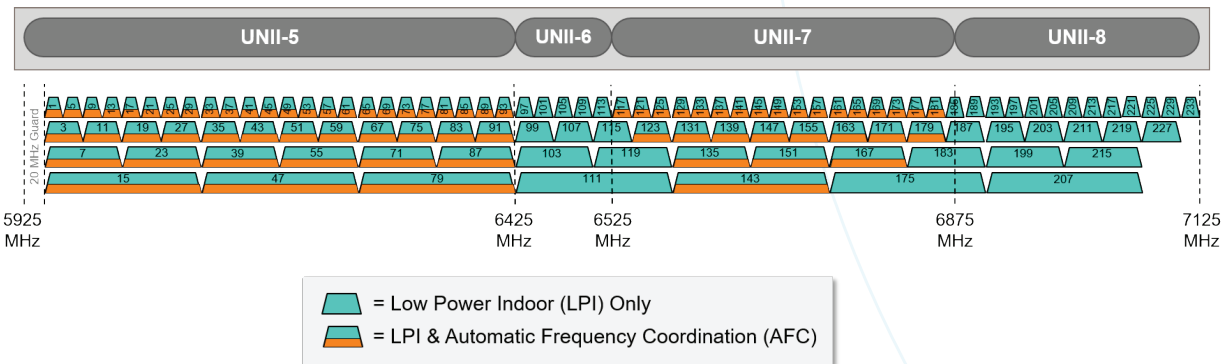


Figure 4: Allowed Outdoor AFC Channels (Orange Stripes)



First, countries that permit AFC do not allow it on all the channels where LPI is permitted. In the U.S., AFC is not allowed on UNII-6 or UNII-8 channels, including those that overlap with UNII-7. This results in the channel plan shown in Figure 4. The channels with orange stripes are those on which AFC is permitted. You can see that the wider the bandwidth in use, the more spectrum is lost due to this constrain.

The second reason is that channel availability is further dependent on what incumbents are operating nearby. In most metro areas in the U.S., there are hundreds of incumbents, meaning that the AFC may report completely different channel availability on a block-by-block basis. Nor is there any guarantee that any channels at all will be available in a given location. Figure 4 simply shows what channels are permitted in the U.S., not what channels are *available* at a given coordinate.

While Aruba expects outdoor 6 GHz networks to be a significant investment catalyst due to the general increase in capacity, the unpredictability for some customers will inevitably lead to indoor and outdoor channel plans being quite distinct. If it becomes desirable to separate outdoor IoT or security video traffic from employee/guest traffic, then it will be necessary to partition the outdoor deployment accordingly. This could be accomplished with co-located outdoor APs to minimize cabling and installation costs.

Note that outdoor stadiums will require APs that support AFC. This applies to the common underseat deployment model where an AP is placed into a protective enclosure. Such APs would have to be AFC-capable, including the ability to self-locate their position via GPS or some other method. By definition, stadium bowl deployments will have a different channel plan than the rest of the facility, but that is no

different than today's dual-band deployments where wireless architects are already trained to think in terms of frequency segmentation.

HOW DO CLIENTS DISCOVER AND ATTACH TO 6 GHZ ONLY PARTITIONS?

To minimize the time required for Wi-Fi devices to discover 6 GHz APs, new features have been added to Wi-Fi 6E to permit APs to advertise the presence of other radios in beacons. Wi-Fi clients do channel scans when they prepare to roam, and this requires the device to stop and listen for a minimum amount of time on each channel. Extending this to 6 GHz without modification would entail unreasonable roaming delays, as much as six additional seconds per roam (59 channels * 100 milliseconds each). Instead, a 6 GHz APs advertises the existence of its 6 GHz radio on its other supported frequency band(s) via a new information element (IE) in the beacon called a Reduced Neighbor Report (RNR). The client can probe or otherwise behave normally but have access to any 6 GHz SSID via the RNR. Aruba is seeing this work extremely well with the our 6 GHz APs.

That's great for a dual-band SSID like the Corp_SSID from Table 2. But what about the Corp_6Only_SSID? How is that to be discovered when there is no equivalent SSID on the other radio(s)?

The answer is that the cross-band advertisement capability of the RNR does not require the SSIDs to be identical. The whole point is to give a roaming client enough information so that it can jump over to 6 GHz and directly initiate a connection for any preconfigured SSID. Figure 5 shows a RNR from a 5 GHz AP beacon that is advertising a 6 GHz radio with an SSID completely different from the one in the beacon.

```

    Tag: Reduced Neighbor Report
      Tag Number: Reduced Neighbor Report (201)
      Tag length: 17
    Neighbor AP Information
      TBTT Information Field Header: 0x0d00
      Operating Class: 131
      Channel Number: 185
      TBTT 0:
        Neighbor AP TBTT Offset: 0x20
        BSSID: ArubaaHe_68:3f:a0 (1c:28:af:68:3f:a0)
        Short SSID: b9577d20
      BSS Parameters: 0x4c
        ... ..0 = OCT Recommended: False
        ... ..0. = Same SSID: False
        ... .1.. = Multiple BSSID: True
        ... .1.. = Transmitted BSSID: True
        ...0 ... = Member of ESS with 2.4/5 GHz Co-Located AP: False
        ..0. .... = Unsolicited Probe Responses Active: False
        .1.. .... = Co-Located AP: True
        0... .... = Reserved: False
      20 MHz PSD: 254
  
```

Figure 5: Reduced Neighbor Report Pointing to 6 GHz BSSID & ESS



The client learns from this IE that there is another radio on this AP and that it is tuned to channel 185 in the 6 GHz band. It also obtains two different methods to initiate an association with that radio. The first method is to simply take the BSSID and directly connect to it (because the 802.11 association process is based on BSSIDs, not the ESSID name). Alternatively, Wi-Fi 6E includes a feature called a “Short SSID” name which is a four-byte hash of the full SSID name. A client that went this route could switch to 6 GHz and do a wildcard probe for all BSSIDs that support the Short SSID name. There are plausible reasons why a client might take either approach. The key point is that neither of these methods requires the same SSID name to exist on multiple radios on the same AP. The only requirement is that there be at least one active SSID on 5 GHz to advertise the RNR. So 6 GHz-only SSIDs are absolutely feasible, and Aruba has successfully tested this configuration.

BAND PARTITIONING IN COUNTRIES OPENING 500 MHZ ONLY

Many countries have taken the bold decision to open all 1,200 MHz of the 6 GHz band for license-exempt Wi-Fi (the “Americas” or FCC model). Yet other countries are taking an incremental approach and opening the lower 500 MHz only (the “European” or ETSI model). Many of these countries may eventually open the upper 700 MHz, but some will not. While this paper focuses on 1,200 MHz partition scenarios to show what is possible, and because these countries constitute a huge percentage of global GDP, the partitioning constructs described can be applied in 500 MHz.

Countries following the European ETSI model also generally do not permit the use of the upper portion of the 5 GHz band including channels 149 – 165 (although this is expected to begin changing soon) or they permit these channels at just 14 dBm. As a result, it may be difficult to subdivide the 5 GHz band but partitioning the 6 GHz band into two 250 MHz blocks is absolutely feasible.

WHAT ABOUT 2.4 GHZ?

The arrival of 6 GHz lays out a clear path to a future where we can finally deprecate 2.4 GHz for normal end-user devices. As explained earlier, the most obvious long-term use of 5 GHz is to provide service to guest, visitor, or transient devices, while 6 GHz is reserved for enterprise devices.

When realized, this means that 2.4 GHz can finally become a pure IoT band. In today’s world, this is a much better application. This can include non-Wi-Fi technologies like Bluetooth and Zigbee, as well as low-data rate IoT devices that use Wi-Fi exclusively. Everyone wins in this approach. General data devices do not have to contend on 2.4 GHz, so one’s Bluetooth headset and mouse will no longer interfere with Zoom calls. And IoT devices benefit from not competing with general Wi-Fi end user traffic.

High data rate IoT devices – most notably HD and UltraHD security cameras – deserve to be put on 6 GHz when support becomes available. The 2.4 GHz band as a rule is for devices with peak burst rates well under 10 Mbps.

BEYOND WI-FI: THE MULTI-RAN ENTERPRISE

The concept of spectrum layering can be extended beyond unlicensed bands and Wi-Fi. In 2021, enterprise wireless architects have more choices for radio access network (RAN) technologies to solve different problems than ever before. Privately-owned 4G or 5G systems operating in the 3.5 GHz Citizens Broadband Radio Service (CBRS) band are real. Aruba offers a CBRS infrastructure solution and there are now over 200 different CBRS-capable endpoint devices on the market. The same logic that was applied above to segmenting the 6 GHz band into two or more partitions can be applied to distribute devices between unlicensed and shared spectrum bands.

WHAT’S NEXT?

The 6 GHz band has enough capacity to provide significant latitude that enterprise wireless architects have never before enjoyed and it allows them to pursue layered approaches that were not feasible in the past. It is not enough to simply cut and paste best practices from 5 GHz to the 6 GHz band. And, while 160 or 80 MHz channels that 6 GHz offers will be needed for certain industries and use cases, consuming that amount of channel bandwidth for IoT or voice devices is potentially very inefficient. With the addition of 6 GHz we move beyond a one-size-fits-all RF planning strategy.



Instead, wireless architects can conceptualize and deploy 6 GHz and/or 5 GHz layers with different channel sizes to support different use cases. For example, a hospital may want to reserve the 6 GHz band for clinical use with multiple layers, each dedicated to a different channel size for a specific use, while allocating the 5 GHz band to guest services. Other industries such as manufacturing and higher education will no doubt create flexible approaches to best meet their needs.

Given adequate spectrum resources, Wi-Fi and other license-exempt technologies will continue to be fertile ground for innovation, coming up with new ways to help organizations of all kinds solve thorny connectivity challenges – just as envisioned by the regulators around the globe who have opened this vital new band. As a result, enterprise wireless architects have more choices for radio access network (RAN) technologies to solve different problems than ever before.