

BEYOND SPEED: TRANSITIONING FROM WI-FI 7'S MULTI-GIGABIT FOUNDATIONS TO WI-FI 8'S ULTRA-RELIABLE NETWORKS - A ROADMAP FOR OPERATORS

A technical paper prepared for presentation at SCTE TechExpo25

Ian Wheelock
CTO, Europe
Vantiva
Ian.Wheelock@vantiva.com

Table of Contents

Title	Page Number
1. Introduction	5
1.1. The evolution of Wi-Fi Design Objectives	5
1.2. Market Evolution and Deployment Realities	6
1.3. Technical Innovation	6
1.4. Residential Deployment Implications and Strategic Considerations	7
1.5. Research Approach	7
2. Standardisation Overview	8
2.1. IEEE 802.11 Working Group	8
2.2. Wi-Fi Alliance	8
2.3. Delivering Wi-Fi Standards and Specifications	9
2.4. Additional Wi-Fi Related Specifications	10
3. Standards Details	10
3.1. IEEE 802.11be - Extremely High Throughput (EHT)	11
3.1.1. Intent and Scope	11
3.1.2. Objectives and Problem Statement	11
3.1.3. Key Technical Directions	11
3.1.4. Dependencies, Process, and Relationship to 802.11ax	11
3.2. IEEE 802.11bn - Ultra High Reliability (UHR)	12
3.2.1. Intent and Scope	12
3.2.1. Objectives and Problem Statement	12
3.2.1. Key Technical Directions	12
3.2.2. Dependencies, Process and Relationship to 802.11be/802.11ax	12
3.3. Summary of Features Across IEEE Releases	13
4. 802.11be/Wi-Fi 7 Features	14
4.1. Overview	14
4.2. 320 MHz Ultra-Wide Channel Operation	14
4.3. Multi-Link Operation (MLO)	15
4.4. 4096- Quadrature Amplitude Modulation (4K-QAM)	16
4.5. Multiple Resource Units (MRU) Allocation	16
4.6. Preamble Puncturing (Flexible Channel Utilization)	17
4.7. BSS Colouring and Spatial Reuse Groups (SRG)	18
4.7.1. BSS Colouring, SRG, and non-SRG - the cocktail party explanation	19
4.8. 802.11be feature mapping to Wi-Fi 7 Release 1	21
4.8.1. Features in 802.11be not in Wi-Fi 7	22
5. 802.11bn/Wi-Fi 8 Features	23
5.1. Overview	23

5.2.	Co-ordinated Spatial Reuse (Co-SR)	24
5.3.	Co-ordinated Beam Forming (Co-BF).....	25
5.4.	Co-ordinated Time Division Multiple Access (Co-TDMA).....	27
5.5.	Coordinated Restricted Target Wake Time (Co-RTWT).....	29
5.6.	Distributed Multi-Link Operation (dMLO) - Seamless Mobility	29
5.7.	Spectrum Utilisation Features.....	31
5.7.1.	Dynamic Sub-Channel Optimisation (DSO).....	31
5.7.2.	Non-Primary Channel Access (NPCA)	32
5.7.3.	Dynamic Bandwidth Expansion (DBE)	33
5.7.4.	Distributed Resource Units (DRU).....	34
5.7.5.	New MCS.....	36
5.7.1.	Unequal Modulation (UEQM).....	37
5.7.2.	2x LDPC	37
5.8.	Enhanced Quality of Service and Reliability	38
5.8.1.	High Priority EDCA	38
5.8.2.	Enhanced Long Range (ELR)	39
5.8.3.	TxOP Pre-emption	40
5.9.	Coexistence Feature.....	41
5.10.	Power Saving.....	42
6.	Upgradeability of Standards	43
6.1.	Backward Compatibility Foundations.....	43
6.2.	Layered Feature Negotiation	44
6.3.	Firmware versus Hardware Evolution.....	44
6.4.	Coexistence of Multiple Generations	45
7.	Summary and conclusion	46
7.1.	“Wi-Fi 7 is the on-ramp; Wi-Fi 8 is the destination” - Executive Summary	46
7.1.1.	Key takeaways.....	47
7.1.2.	Call to action.....	47
7.2.	What Wi-Fi 7 Delivers Today (foundation layer).....	47
7.2.1.	Practical capabilities in deployed gateways/extenders.....	47
7.2.2.	Deployment reality	47
7.2.3.	Why this foundation matters for Wi-Fi 8.....	48
7.3.	Wi-Fi 8: Feature-to-Outcome Mapping	48
7.4.	Why Multi-AP/Mesh Is Required for Wi-Fi 8 Value Realization	49
7.5.	Operator Rollout Playbook.....	49
7.5.1.	Mesh acceleration options	49
7.5.2.	Competitive positioning.....	49
7.5.3.	Business framing (directional, no hard costs).....	49
7.5.4.	Consumer Value Narrative	49
7.5.5.	Risks, Dependencies, and Mitigations.....	50
7.5.6.	Metrics and Proof Points.....	50
7.5.7.	FAQ for Executives and Care Agents.....	51
7.6.	Recommendation.....	51
	Abbreviations.....	51
	Bibliography & References	54

List of Figures

Title	Page Number
Figure 1, IEEE TG and Wi-Fi Alliance Timelines	9
Figure 2, Wi-Fi 4 to Wi-Fi 8 Feature/Capability Additions.....	14
Figure 3, 320 MHz Ultra-Wide Channels in 6 GHz.....	15
Figure 4, Multi-Link Operation - EMLMR/EMLSR.....	16
Figure 5, 1024-QAM and 4096-QAM.....	16
Figure 6, Multiple Resource Units (MRU).....	17
Figure 7, Multiple-RU and Preamble Puncturing	18
Figure 8, Basic BSS Colouring, Three APs, Three Different Colours.....	19
Figure 9, Complex BSS Colouring with SRG and Non-SRG groups.....	21
Figure 10, Co-ordinated Spatial Reuse (Co-SR)	25
Figure 11, Co-ordinated Beam Forming (Co-BF)	26
Figure 12, Co-ordinated TDMA	28
Figure 13, MLO and distributed-MLO Operation (Seamless Mobility Domain)	30
Figure 14, DSO sub-channel options	32
Figure 15, NPCA Channel Switching Options	33
Figure 16, Dynamic Bandwidth Expansion (DBE).....	34
Figure 17, RRU/MRU versus Distributed RU (DRU)	35
Figure 18, Maximum Transmit Power of RU (dBm).....	35
Figure 19, New MCS (17, 19, 20 & 23)	36
Figure 20, TxOP Preemption in action	41
Figure 21, Wi-Fi 6/7 to Wi-Fi 8 Adoption Options.....	46

List of Tables

Title	Page Number
Table 1, 802.11 Task Group Timelines	10
Table 2, Wi-Fi Alliance Wi-Fi Release Timelines.....	10
Table 3, IEEE Wi-Fi Version Features Compared.....	13
Table 4, 802.11be Features in Wi-Fi 7 Release 1 and 2.....	22
Table 5, 802.11be Features not included in Wi-Fi CERTIFIED 7.....	23
Table 6, Legacy Wi-Fi devices vs Wi-Fi 8	48

1. Introduction

As the demand for faster and more reliable wireless networks continues to grow, the transition from Wi-Fi 7 to Wi-Fi 8 represents a significant leap forward in technology. While Wi-Fi 7 has laid the groundwork with its multi-gigabit speeds, Wi-Fi 8 aims to deliver ultra-reliable networks that cater to the evolving needs of modern applications. This document provides a comprehensive roadmap for operators to navigate this transition, highlighting the key advancements and benefits of Wi-Fi 8.

Wi-Fi 8 introduces several groundbreaking features designed to enhance network performance, reliability, and efficiency. Among these innovations is Coordinated Time Division Multiple Access (Co-TDMA), which replaces the traditional contention-based access methods with a deterministic scheduling approach, along with Coordinated Spatial Reuse and Coordinated Beamforming. This shift promises to reduce latency, improve channel efficiency, and support the stringent requirements of real-time applications such as augmented reality (AR), virtual reality (VR), and industrial automation.

The journey to Wi-Fi 8 is not just about speed; it's about creating networks that can deliver consistent and predictable performance in increasingly dense environments. By understanding the technical advancements and strategic considerations outlined in this document, operators can position themselves at the forefront of this technological evolution, ensuring they meet the demands of their customers and stay ahead of the competition.

Wi-Fi 7 is aligned to the IEEE 802.11be standard, while Wi-Fi 8 builds on the features being developed through the IEEE 802.11bn task group work, and will then be curated into the Wi-Fi Alliance's Wi-Fi Certified 8 certification program. 802.11bn delivered their Draft 1.0 release at end of July 2025¹, with a final report due in 2028, while Wi-Fi Alliance began Wi-Fi 8 work in May 2025, with a target launch of certification by end 2027/start 2028.

1.1. The evolution of Wi-Fi Design Objectives

The Wi-Fi networking landscape has undergone significant transformation over the past 10-20 years, driven largely by technological competition that has seen enormous increases in bandwidth delivery to homes and businesses. Back in 2000, a DOCSIS 1.0 system shared approximately 30 Mbps between 1,000 connected homes - today, 10,000 Mbps is shared between 100-200 homes. The addition of Wi-Fi to the connected home drove the need to match the WAN speeds. The evolution of Wi-Fi tried to keep up with these major WAN speed upgrades and the increasing demands of streaming media and gaming applications, maximising throughput while maintaining cost-effective solutions and often hitting the limit of hardware capability at the time. Wi-Fi 7, with its focus on multi-gigabit speeds and ultra-low latency, represents the peak of the throughput-centred approach. However, as Wi-Fi networks continue to densify to provide maximum throughput and coverage, and application requirements become more diverse, it has been recognised that raw speed alone is not sufficient for modern Wi-Fi deployments. The densification also leads to challenges where multiple APs are competing against each other for Wi-Fi access and doing so inefficiently.

The average home in the US has between 20 and 25 devices²³, including smartphones, laptops, smart thermostats, security cameras, and IoT sensors. This huge increase in devices and the growing demand for bandwidth-intensive applications such as 4K streaming, cloud gaming, and video conferencing has resulted in new challenges that cannot be fixed through speed alone.

Wi-Fi 8 has emerged as a response to the needs of the evolving network and application demands, introducing a reliability-centred design approach that focuses on network coordination, predictable

performance, and energy efficiency. This shift reflects a better understanding of real-world deployment challenges, where interference, coverage gaps, mobility, and battery life constraints matter as much as peak throughput. Unlike previous Wi-Fi generations, the usual tools of boosting the MCS by 20% or doubling the Wi-Fi channel bandwidth (due to new spectrum in 6 GHz) would not provide practical benefits (for example, 8192-QAM would offer increased throughput but range would be significantly reduced, and equally, 640 MHz channels in 6 GHz spectrum would not have been applicable in regions with only 480 MHz of spectrum)

1.2. Market Evolution and Deployment Realities

The transition from Wi-Fi 6 to Wi-Fi 7 has provided solid insight into the realities of next-generation residential Wi-Fi deployments. Headline Wi-Fi 7 features such as 320 MHz channels, 4096-QAM, and Multi-Link Operation (MLO) promise impressive gains, but operators and homeowners have learned that extracting consistent benefits in dense, interference-prone environments requires careful RF planning and more sophisticated Wi-Fi management. MLO, in particular, introduces significant advantages; however, early deployments show that multi-vendor interoperability and roaming behaviours are still maturing.

Single-family homes, especially larger, multi-floor dwellings, have identified the limits of single access point (AP) deployments and the need for multi-AP mesh to deliver stable, whole-home coverage. Multi-dwelling units (MDUs) pose even more challenges: dozens of neighbouring APs create heavy co-channel and adjacent-channel interference that can constrain performance regardless of the Wi-Fi generation in use.

Many Wi-Fi 7 tri-band routers and mesh systems support 320 MHz channels on 6 GHz, but practical use depends on client capability, regional spectrum rules, and short-range conditions. 4096-QAM is highly range and SNR-dependent and degrades rapidly through attenuating materials such as walls, doors, and windows. While some high-end laptops and flagship phones/tablets support 320 MHz at 6 GHz, the majority of clients today still operate at 80 or 160 MHz. In mesh designs, 6 GHz extenders/satellites that support 320 MHz can provide improved dedicated backhaul between nodes, helping maximize end-to-end throughput across the home.

Wi-Fi 8's coordinated approaches directly address these residential challenges, introducing network-level features that allow APs to work together to minimise interference, optimize channel and airtime utilization, and deliver more consistent performance across mixed client populations. This evolution shifts the focus from peak throughput to reliability and predictability, with meaningful implications for how residential Wi-Fi is planned, deployed, and operated.

It's also important to recognize the side-benefits of Wi-Fi 8. The primary gains apply to Wi-Fi 8 capable devices (e.g., Wi-Fi 8 extenders or clients), while coordination and efficiency improvements can indirectly help legacy devices by reducing contention and improving airtime availability. As with previous generations, initial Wi-Fi 8 rollouts will be single-AP or multi-AP mesh. A Wi-Fi 8 mesh can immediately leverage coordination features across the APs themselves, improving the network's RF behaviour for all devices using the same spectrum, even if many clients are legacy. As customer devices turn over (upgrades/replacements), more traffic moves onto Wi-Fi 8 MAC/PHY features, netting direct improvements for those clients and freeing up airtime for older devices sharing the AP.

1.3. Technical Innovation

Wi-Fi 8 builds on Wi-Fi 7's feature set while adding coordination mechanisms that amplify the value of those features in real homes. For example, Wi-Fi 7 introduced Multi-Link Operation (MLO) as a device-

centric capability; Wi-Fi 8 evolves this into Distributed MLO (D-MLO), enabling multi-AP, multi-band coordination and more seamless mobility across mesh nodes.

Similarly, Wi-Fi 8's Coordinated Target Wake Time (TWT) enhances Wi-Fi 7's power-save capabilities into a network-wide scheduling system. In residential settings, this enables better orchestration of activity across user devices and IoT endpoints (e.g., cameras, sensors, locks), reducing contention and improving battery life. In larger enterprise environments, the same concept scales to synchronize access across thousands of devices to maximize predictability and energy efficiency.

Multi-AP Coordination in Wi-Fi 8 is especially relevant for residential mesh. By sharing real-time information-channel conditions, interference levels, client requirements-mesh nodes can coordinate channel access, spatial reuse, and scheduling. Instead of competing, nodes collaborate to reduce collisions and retries, balancing load and making far more efficient use of limited spectrum.

1.4. Residential Deployment Implications and Strategic Considerations

Wi-Fi 7 delivered immediate improvements in throughput and efficiency that improve conditions for streaming, gaming, and work-from-home-particularly in multi-user households. These gains can be limited, however, by the WAN speed tier (e.g., sub-gigabit internet) or by older client devices that cannot take advantage of newer Wi-Fi 7 features.

The Wi-Fi 7 to Wi-Fi 8 upgrade changes the residential value proposition. Rather than emphasizing only higher peak rates, Wi-Fi 8 focuses on reliability, determinism, and coordination. Homes with extensive smart-home deployments, challenging RF conditions (e.g., MDUs, dense neighbourhoods, heavy multipath), or requirements for consistent performance (e.g., video calls, AR/VR, security feeds) stand to benefit significantly from Wi-Fi 8's coordination features (assuming these devices migrate to Wi-Fi 8). Energy-saving improvements via Coordinated rTWT will be very appealing to households with dozens of battery-powered devices, both extending lifetime and reducing maintenance effort, but again requiring device replacement.

As with every generation shift, upgrading the APs typically yields some benefits for legacy clients thanks to better radios, CPUs/NPUs, and scheduling algorithms; the largest gains occur when clients and mesh nodes also match the AP Wi-Fi generation. In particular, adding same-gen mesh nodes can immediately improve backhaul capacity and stability (e.g., 6 GHz / 320 MHz), which lifts performance for all clients. Even when only some clients use advanced features (e.g., Wi-Fi 7's 320 MHz or 4K-QAM under ideal conditions), their reduced airtime consumption frees capacity for legacy devices. For example, if a Wi-Fi 7 client completes transmissions in 30% less airtime due to wider channels or higher-order modulation, that airtime becomes available for other devices-improving overall network responsiveness.

1.5. Research Approach

This paper draws from multiple sources, including IEEE 802.11 working group materials, Wi-Fi Alliance technical specifications, consumer router vendor roadmaps, residential ISP deployment data, and testing results from Wi-Fi 7 and Wi-Fi 8 implementations in home environments. Given the evolving nature of the Wi-Fi 8 (802.11bn) standard, much of the technical insight reflects current IEEE TGBn Ultra High Reliability (UHR) draft work rather than finalized specifications, with Draft 1.0 completed and final ratification expected in 2028.

The paper prioritises features and capabilities that have a direct impact on residential deployments, focusing on mesh networking architectures, smart-home integration, and mixed-device environments common to single-family units (SFUs) and multi-dwelling units (MDUs). While some enterprise concepts

are referenced for context, the main focus is on residential use-cases where reliability, coverage, and ease of use are most important.

Residential deployments have specific requirements: multi-AP mesh for coverage, co-existence of legacy and new clients, backward compatibility, and user-friendly operation within cost constraints. The paper also considers economic practicality and long-term value for both households and service providers.

The document is structured as follows: an introduction to Wi-Fi 7 and Wi-Fi 8 design objectives in residential contexts; a comparative analysis of key differences between EHT (802.11be) and UHR (802.11bn); upgradeability paths across generations; adoption timelines and availability; and a concluding overview on implications for the evolution of home networking

2. Standardisation Overview

The development and certification of Wi-Fi standards involve a coordinated effort between the IEEE 802.11 Working Group and the Wi-Fi Alliance. Understanding the roles and processes of these organisations is essential for service providers, consumer electronic manufacturers, and mesh system vendors.

2.1. IEEE 802.11 Working Group

The IEEE 802.11 Working Group is responsible for developing and maintaining the technical standards for wireless local area networks (WLANs). The process begins with a Project Authorization Request (PAR) that defines the scope, objectives, and timeline of a new standard. The group includes various bodies such as Study Groups (SG), Task Groups (TG), and Standing Committees (SC). The development process involves identifying needs, creating and approving the PAR, forming TGs for technical development, and final balloting and publication of the standard.

For example, Wi-Fi 7 (802.11be) was authorised in May 2019 to build upon Wi-Fi 6's efficiency and support "Extremely High Throughput" (EHT). Its goals included achieving multi-gigabit throughput, backward compatibility, and improved latency and jitter. These objectives translated into technological advancements like Multi-Link Operation, wider 320 MHz channels, and 4096-QAM modulation.

More recently, Wi-Fi 8 (802.11bn), approved in September 2023, focuses on "Ultra High Reliability" (UHR). Its objectives include increasing throughput, reducing latency and data loss, lowering power consumption, and improving peer-to-peer operations, while maintaining backward compatibility.

2.2. Wi-Fi Alliance

The Wi-Fi Alliance is a global non-profit organisation that manages certification programs to ensure Wi-Fi devices from different manufacturers work seamlessly together. The Alliance creates certification programs, technical specifications, and marketing campaigns to promote Wi-Fi technology. The certification process begins with identifying market needs and garnering interest from member companies. A Marketing Task Group (MTG) produces a Marketing Requirements Document (MRD), and a Technical Task Group (TTG) develops detailed technical requirements.

The TTG's work covers the creation of the formal specification and development of comprehensive test plans. These test plans, once finalised are used by Approved Test Labs (ATLs) to conduct the Wi-Fi Alliance certification testing to ensure devices meet interoperability and security criteria.

2.3. Delivering Wi-Fi Standards and Specifications

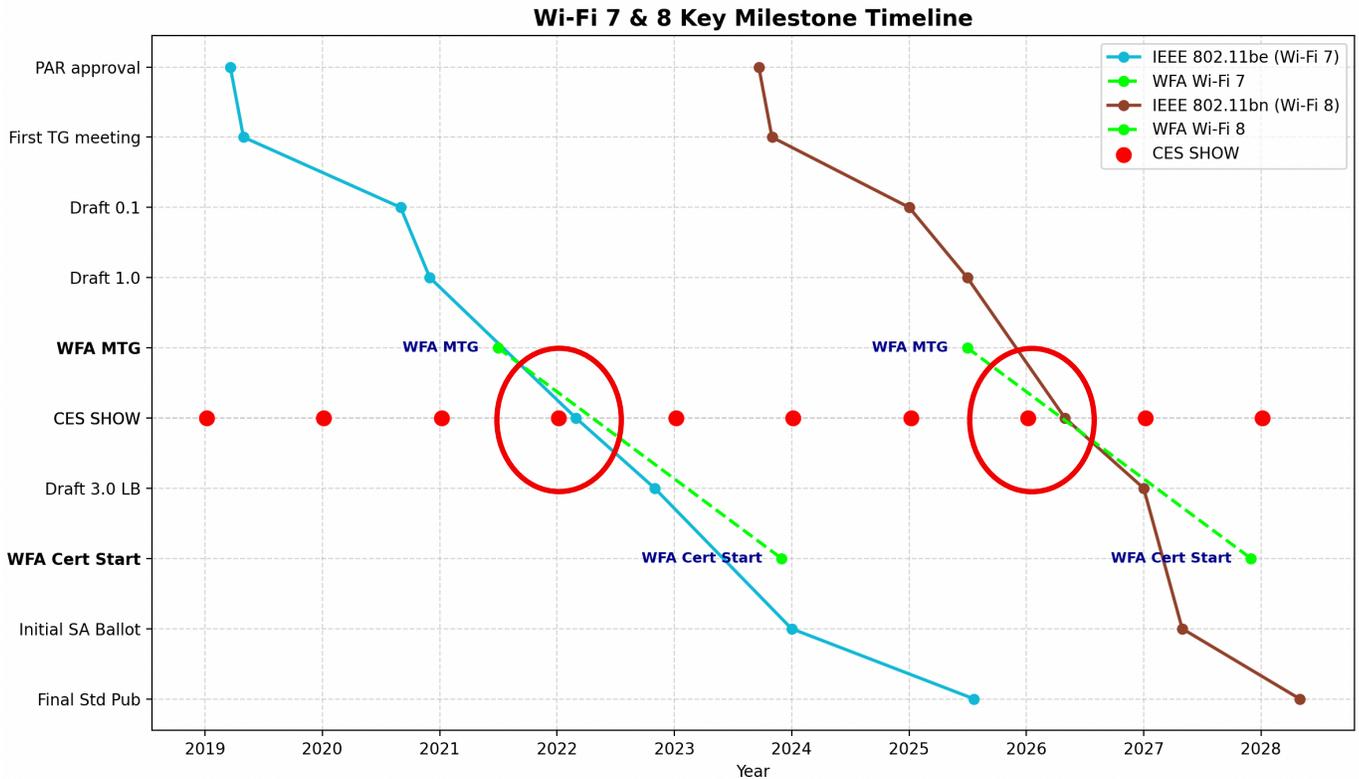


Figure 1, IEEE TG and Wi-Fi Alliance Timelines

This graphic shows the overlap of IEEE and WFA activities in terms of getting to a certifiable Wi-Fi solution. The two IEEE lines show the reality (802.11be) of time taken to get the specification completed through the various approval processes of the IEEE Task Groups. The chart also shows when the Consumer Electronics Show (CES) occurs, where a number of Wi-Fi suppliers tend to get their early versions of the developing specifications into the public domain, showing off their engineering capability, and demonstrate some of the fundamental technology being developed as part of the latest Wi-Fi release. With IEEE 802.11bn underway since September 2023, it's highly likely there will be some display of new Wi-Fi solutions targeting Wi-Fi 8 at CES 2026. For example (shown in red circle above), MediaTek demonstrated Wi-Fi 7 technology at CES2022, 6 months after the Wi-Fi Alliance MTG formation, and nearly two years before the Wi-Fi 7 Certified program started certifying devices. Some silicon companies have already declared that they will be showing off some capabilities of Wi-Fi 8 at CES2026.

The following table identifies the historical timelines for Wi-Fi 7 and some of the historical and rough timelines for delivering Wi-Fi 8.

Event	IEEE 802.11be (Wi-Fi 7)	IEEE 802.11bn (Wi-Fi 8)
PAR approval	Mar-2019	Sep-2023
First TG meeting	May-2019	Nov-2023
Draft 0.1	Sep-2020	Jan-2025
Draft 1.0	Dec-2020	Jul-2025
Draft 2.0 WG Letter Ballot	Mar-2022	May-2026
Draft 3.0 LB	Nov-2022	Jan-2027
Initial SA Ballot	Jan-2024	May-2027
Final Std Pub	Jul-2025	May-2028

Table 1, 802.11 Task Group Timelines

Event	Wi-Fi 7 R1	Wi-Fi 8
WFA MTG SoW	Jul-2021	Jul-2025
WFA Cert Start	Dec-2023	Dec-2027

Table 2, Wi-Fi Alliance Wi-Fi Release Timelines

2.4. Additional Wi-Fi Related Specifications

In addition to the main Wi-Fi standards, several other specifications support the operation and management of Wi-Fi systems. New features defined by Wi-Fi 8 will eventually be supported in these specifications (where appropriate). These include:

- Wi-Fi Protected Access (WPA) 3: for secure operation of the Wi-Fi network.
- IEEE 802.11k: Radio Resource Management information.
- IEEE 802.11r: Fast BSS Transition - Fast Roaming.
- IEEE 802.11v: Wireless Network Management Steering/Optimization.
- Wi-Fi EasyMesh: Multi-AP coordination and whole home Wi-Fi.
- Wi-Fi Data Elements: Standardised Wi-Fi Telemetry.
- Wi-Fi Sensing: Motion/Awareness via Wi-Fi Signals.
- Hotspot 2.0/Passpoint: Carrier-grade Onboarding and Roaming.

3. Standards Details

This section introduces the high level scope, as defined by the IEEE PAR for 802.11be and 802.11bn specifications. The primary focus on 802.11be is “**Extremely High Throughput (EHT)**” and it has capitalised on the use of 320 MHz/4096- /4 096 QAM operation, primarily in the 6 GHz spectrum and has coupled that with Multi Link Operation (MLO) to augment the maximum possible speeds achievable with Wi-Fi. 802.11bn is more focussed on “**Ultra High Reliability (UHR)**” rather than throughput expansion. The features in 802.11bn enable the more effective use of the spectrum through a number of co-ordination features (beamforming, spatial reuse, TDMA) as well as optimising use of the existing channel

bandwidth. Range extension is supported, particularly for FCC based power regulations through features like Distributed Resource Units (DRU), while mobility and roaming between APs in a multi-AP network is also focussed on, extending the operation of MLO across multiple APs. The next sections discuss the background and “reason for” each of these two specifications, followed by a deeper dive in to the Wi-Fi 7 and Wi-Fi 8 features.

3.1. IEEE 802.11be - Extremely High Throughput (EHT)

3.1.1. Intent and Scope

802.11be is an amendment to IEEE 802.11 aimed at “Enhancements for Extremely High Throughput (EHT)”. The amendment made modifications to both the PHY and MAC to enable at least one mode supporting a maximum throughput of at least 30 Gbps at the MAC Service Access Point (SAP), with carrier frequency operation between 1-7.250 GHz, while ensuring backward compatibility and coexistence with legacy 802.11 devices in 2.4 GHz, 5 GHz, and 6 GHz bands. The amendment also defined at least one mode with improved worstcase latency and jitter.

3.1.2. Objectives and Problem Statement

Meet the evolving needs of high-bandwidth and low-latency applications such as 4K/8K video, AR/VR, gaming, remote office/cloud computing, with expectations like <5ms latency for real-time gaming.

Improve aggregate throughput, latency, reliability, and power efficiency to maintain WLAN competitiveness and better integrate with Time Sensitive Networking (TSN) in heterogeneous Ethernet/WLAN environments. Continued co-existence with previous 802.11 versions (e.g. 802.11ax)

3.1.3. Key Technical Directions

Candidate areas included:

- 320 MHz bandwidth and more efficient use of non-contiguous spectrum
- Multi-band/multi-channel aggregation and operation
- Up to 16 spatial streams and MIMO protocol enhancements
- Enhanced link adaptation and retransmission (e.g., HARQ)
- Adaptation to 6 GHz regulatory rules as needed

Early TGbe discussions explored multi-AP ideas (e.g., coordinated/joint transmission), but the task group deferred the concept of a formal multi-AP coordination framework due to timeline/complexity, leaving it for TGbn.

3.1.4. Dependencies, Process, and Relationship to 802.11ax

The PAR indicated that completion of 802.11be depended on P802.11ax, as 802.11be builds on the High Efficiency (HE) foundation established by 802.11ax. It extended OFDMA, spatial reuse/efficiency methods, and 6 GHz operation with Enhanced High Throughput (EHT) features such as multi-link operation (MLO), 320 MHz channels, higher modulation schemes, deeper aggregation, and larger MIMO configurations. These enhancements were designed to meet substantially higher peak throughput targets and deliver improved latency, jitter, and reliability.

The PAR was approved in March 2019, and the official specification was released in July 2025

3.2. IEEE 802.11bn - Ultra High Reliability (UHR)

3.2.1. Intent and Scope

802.11bn is an amendment that adds an “Ultra High Reliability” capability to WLAN through modifications to both PHY and MAC, operating between 1-7.250 GHz, ensuring backward compatibility and coexistence in 2.4 GHz, 5 GHz, and 6 GHz unlicensed bands. Note that millimetre-wave operation was initially hinted at for 802.11bn, but this is out of scope. Another task group, 802.11bq however is tasked with addressing millimetre-wave operation.

802.11bn focuses on improving throughput, latency, and MPDU loss metrics versus the 802.11be (EHT) baseline and includes mechanisms to reduce AP power consumption and improve peer-to-peer (P2P) operation.

3.2.1. Objectives and Problem Statement

At least one mode capable of increasing throughput by 25% at the MAC SAP in at least one SINR level (rate versus range), compared to EHT, reducing latency by 25% for the 95th percentile of the latency distribution versus EHT, and reducing MPDU loss compared to EHT for a given scenario, especially during transitions between BSSs.

3.2.1. Key Technical Directions

Better coordination between neighbouring APs and between P2P networks due to increased P2P usage sharing medium resources. The concept of coordination rather than competition between APs is a key theme of 802.11bn, where there are multiple benefits from improved beamforming, spatial reuse and even bandwidth sharing when channels are unused. Increased transmit power for stations is also possible, particularly for devices subject to FCC regulations. Extending the operation of MLO to allow for improved roaming between APs, and reducing power consumption for APs (including mobile APs), with attention to energy efficiency and regulatory alignment. Ensure backward compatibility/coexistence across previous Wi-Fi generations.

3.2.2. Dependencies, Process and Relationship to 802.11be/802.11ax

The PAR for 802.11bn identifies its dependency on P802.11be and positions itself as an evolution of the EHT framework. 802.11bn does not increase Wi-Fi throughput via new spectrum or wider channels but instead focuses on advancing existing capabilities. It builds on 802.11be’s EHT features - including MLO, 320 MHz channels, higher-order modulation, and deeper frame aggregation - which themselves were introduced or refined from 802.11ax’s HE features such as OFDMA, BSS colouring, spatial reuse groups, and 6 GHz operation. This progression reflects a clear technical alignment: 802.11ax improved spectrum efficiency and multi-user performance; 802.11be expanded capacity, throughput, and multi-link flexibility; and 802.11bn shifts emphasis toward predictable latency, high-reliability under sustained load, improved peer-to-peer coexistence in dense environments, better roaming and greater energy efficiency. By leveraging established HE and EHT mechanisms, 802.11bn delivers backwards-compatible improvements while delivering better performance and reliability for next-generation WLAN deployments.

The PAR was approved in September 2023, with Draft 1.0 approved in July 2025.

3.3. Summary of Features Across IEEE Releases

The following table describes how features have evolved across the last number of IEEE 802.11 releases. As can be seen the focus on 802.11be built heavily on options provided in 802.11ax, especially on the introduction of 6 GHz capability.

802.11bn does not have the same scope for increasing bandwidth like 802.11be did, but its turning its attention to using the available bandwidth more efficiently

Feature	802.11ax (HE)	802.11be (EHT)	802.11bn (UHR)
OFDMA	Introduced	Enhanced scheduling uplink/downlink symmetry	Further optimized for latency and reliability
BSS Colouring	Introduced	Refined	Optimized for dense deployment performance
Spatial Reuse Groups	Introduced	Refined and more dynamic	Optimized for reliability and spectrum reuse
6 GHz Operation	Introduced (HE in 6 GHz)	Extended for EHT 6 GHz	Maintained tuned for reliability KPIs
Multi-Link Operation (MLO)	Not present	Introduced with up to 3 links	Enhanced link management for reliability
Channel Bandwidth	Up to 160 MHz	Up to 320 MHz	No change (focus on efficiency not BW)
Modulation	Up to 1024-QAM	Up to 4096-QAM	No change in maximum modulation, tuned for stability
Frame Aggregation	Standard HE aggregation	Deeper aggregation limits	Tuned for latency-sensitive flows
MIMO Configurations	Up to 8x8 MU-MIMO	Up to 16x16 MIMO	Tuned for consistent performance
Latency / Reliability Focus	Initial QoS and scheduling improvements	Lower latency/jitter targets	Primary design focus (predictable low latency, low jitter)
Energy Efficiency	Basic power save mechanisms	Improved power efficiency	Major focus with adaptive power usage
P2P Coexistence	Basic coexistence	Improved coexistence in dense environments	Advanced coexistence mechanisms
4096-QAM	Not present	Introduced	Maintained stability-focused
Up to 320 MHz Channels	Not present	Introduced	Maintained efficiency-focused
Multi-AP Coordination	Not present	Mentioned	Expanded coordination scope for reliability
Improved Roaming	Not present	Not present	Distributed MLO allowing seamless roaming between APs

Table 3, IEEE Wi-Fi Version Features Compared

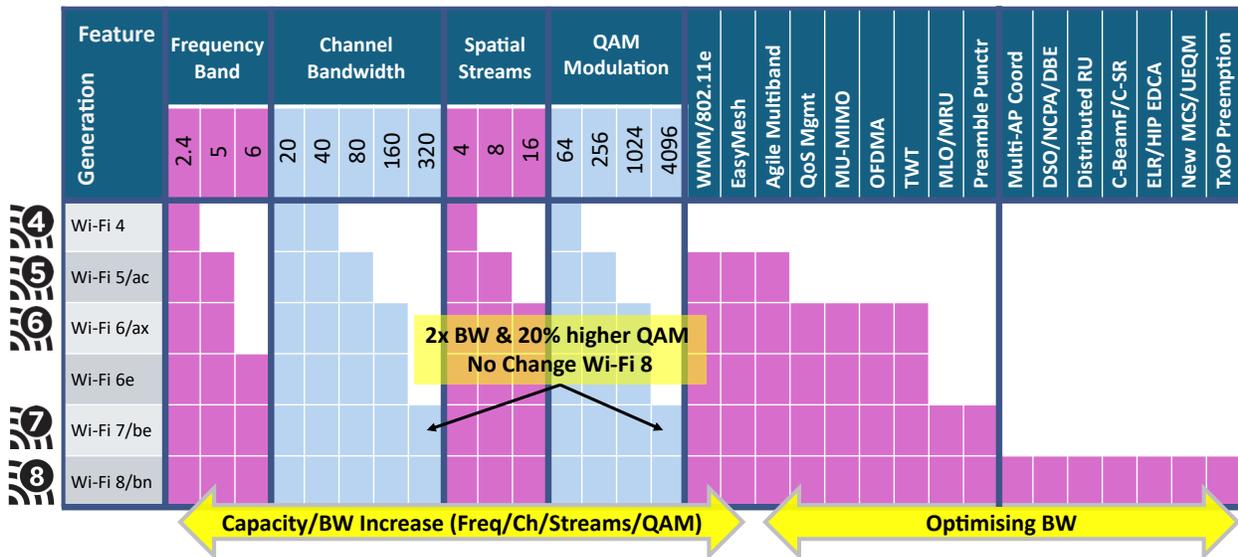


Figure 2, Wi-Fi 4 to Wi-Fi 8 Feature/Capability Additions

4. 802.11be/Wi-Fi 7 Features

4.1. Overview

“Extremely High Throughput” is the name of the game for Wi-Fi 7. The maximum capacity possible with Wi-Fi 7, considering 320 MHz, 4096-QAM and up to 16 spatial streams is a little over 46 Gbps. This is over 4,000 times as much capacity as 802.11b provided when it was introduced in 1999. The maximum speed normally associated with Wi-Fi 7 in a Tri Band configuration (2.4, 5 & 6 GHz) is about 18.6 Gbps (rounded up to BE19000), with all bands operating 4 spatial streams each at maximum channel bandwidth. The following sections detail the key features of Wi-Fi 7, covering throughput, channel aggregation, channel optimization and latency.

4.2. 320 MHz Ultra-Wide Channel Operation

Wi-Fi 7's most transformative feature is the addition of 320 MHz channel bandwidth, a doubling of Wi-Fi 6's maximum 160 MHz channels. This ultrawide bandwidth operates exclusively in the 6 GHz band introduced in Wi-Fi 6e, where up to three 320 MHz channels are available, effectively doubling the data transmission capacity compared to previous generations. Note that recent efforts to get the 7.125 to 7.250 GHz slice of spectrum allocated to Wi-Fi would enable up to 4 ultra-wide channels in 6 GHz. The 802.11be PAR already targeted operation up to 7.250 GHz, now it's down to regulatory efforts.

The 320 MHz channels utilise two 160 MHz carriers based on the Wi-Fi 6 architecture, containing up to 3,920 valid subcarriers compared to Wi-Fi 6's maximum of 1,960 subcarriers in a single transmission instance. This is a huge shift in spectrum utilization, enabling maximum data rates of up to 2.9 Gbps per spatial stream at 320 MHz bandwidth (dependent on the other key Wi-Fi 7 features - 4096-QAM)

Contiguous 320 MHz channels and non-contiguous 160+160 MHz configurations are supported, enabling some flexibility in spectrum allocation addressing both regulatory and interference conditions. Unfortunately, availability of 320 MHz channels is geographically limited due to varying 6 GHz band regulations across different countries, with some regions only offering the lower 500 MHz portion of the

6 GHz spectrum. Other countries, such as China, have no support for Wi-Fi in the 6 GHz spectrum, instead allocating it completely to licensed cellular use. Similarly, ongoing lobbying by cellular communication groups continues to put pressure on IMT use in 6 GHz, at the cost of unlicensed Wi-Fi.

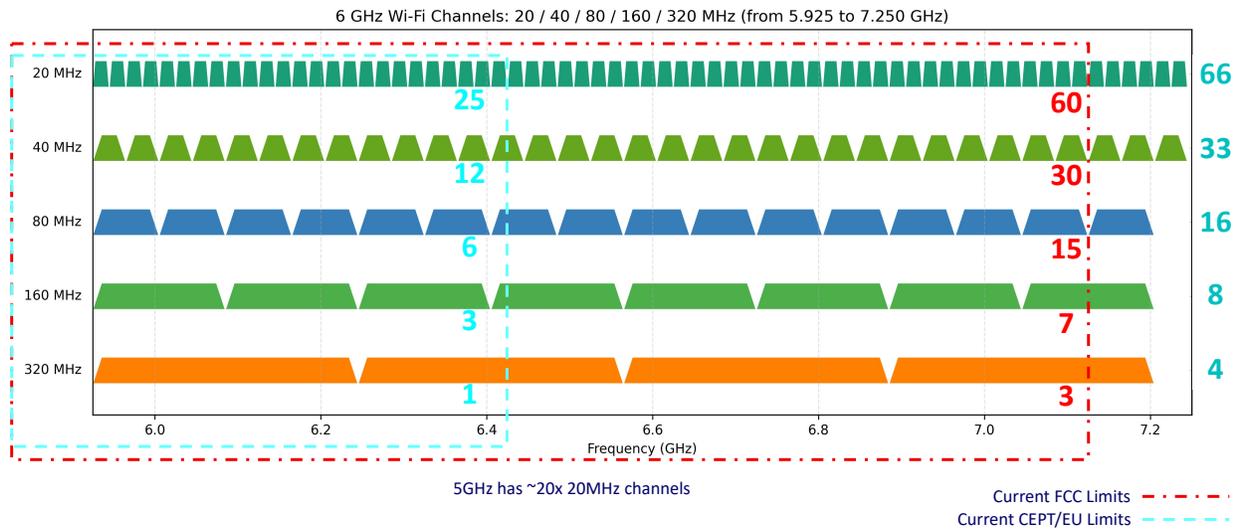


Figure 3, 320 MHz Ultra-Wide Channels in 6 GHz

4.3. Multi-Link Operation (MLO)

Multi-Link Operation represents Wi-Fi 7's biggest innovation, changing how wireless devices establish and maintain network connections. MLO enables Multi-Link Devices (MLDs) to simultaneously establish and operate multiple links concurrently across different frequency bands (2.4 GHz, 5 GHz, and 6 GHz), enabling major boosts in total bandwidth between devices compared to the single-band limit of previous Wi-Fi generations.

The MLO framework supports several operational modes: Simultaneous Transmit and Receive (STR) mode, where multiple links operate independently without mutual interference; Non-Simultaneous Transmit and Receive (NSTR) mode, where links coordinate transmissions and receptions to avoid overlap; Enhanced Multi-Link Single Radio (EMLSR) mode, where a single radio can monitor multiple bands while transmitting on one; and Enhanced Multi-Link Multi-Radio (EMLMR) mode, where multiple radios provide full-duplex, multi-band flexibility with independent channel access per radio.

By enabling traffic load balancing, redundant path reliability, and aggregated bandwidth utilization, MLO delivers significant performance gains. In real-world conditions, it can reduce latency by up to 25% compared to single-link operation⁶, while enabling seamless failover when one link experiences degradation. These benefits make MLO particularly valuable for latency-sensitive use cases such as augmented reality (AR), virtual reality (VR), and real-time gaming, where both throughput and reliability are critical.

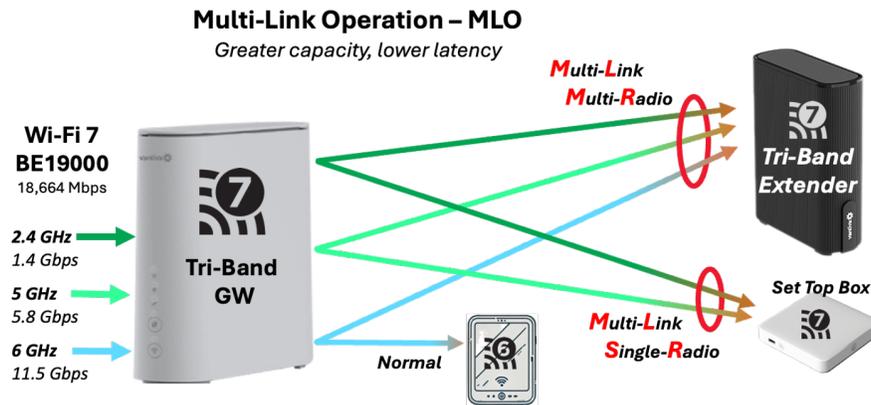


Figure 4, Multi-Link Operation - EMLMR/EMLSR

4.4. 4096- Quadrature Amplitude Modulation (4K-QAM)

4096-QAM modulation provides a 20% improvement in spectral efficiency compared to Wi-Fi 6's 1024-QAM. This higher-order modulation scheme enables each symbol to carry 12 bits of information compared to Wi-Fi 6's 10 bits per symbol, directly translating to increased data transmission rates without requiring additional spectrum resources

This feature however requires significantly higher signal-to-noise ratio (SNR) conditions compared to lower-order modulation schemes, requiring close proximity between transmitter and receiver for optimal performance, primarily in-room and close to AP. The practical deployment of 4096-QAM is expected to be mostly beneficial in scenarios requiring high throughput over short distances, such as enterprise conference rooms, high-density residential environments, and industrial IoT applications

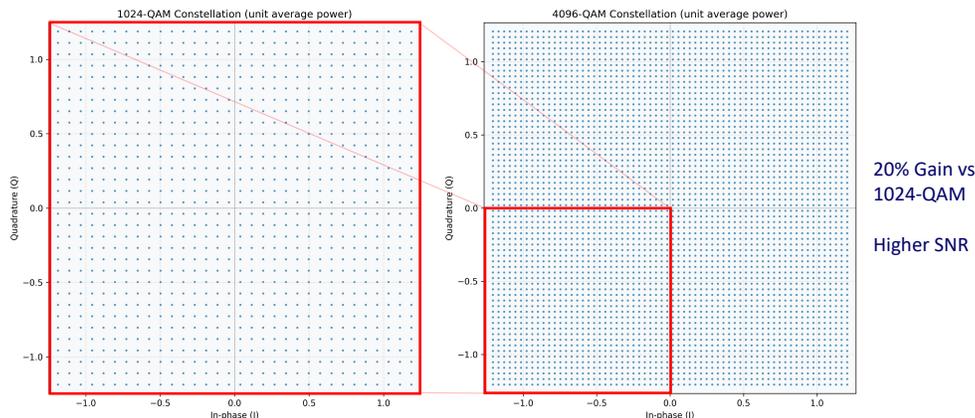


Figure 5, 1024-QAM and 4096-QAM

4.5. Multiple Resource Units (MRU) Allocation

The Multiple Resource Unit (RU) feature represents a major enhancement to Wi-Fi 6's OFDMA technology, addressing the limits of a single RU allocation per user. MRU enables access points to allocate multiple non-contiguous resource units to a single user, significantly improving spectrum utilization efficiency and reducing transmission latency. The MRU implementation allows for dynamic allocation of RU combinations (tones/sub-carrier groups), with some specific restrictions that simplify the implementation while maintaining spectrum efficiency. Small RUs (containing fewer than 242 tones)

can only be combined with other small RUs, while large RUs (242 tones or greater) can only be combined with other large RUs, preventing mixed-size combinations that would complicate receiver processing.

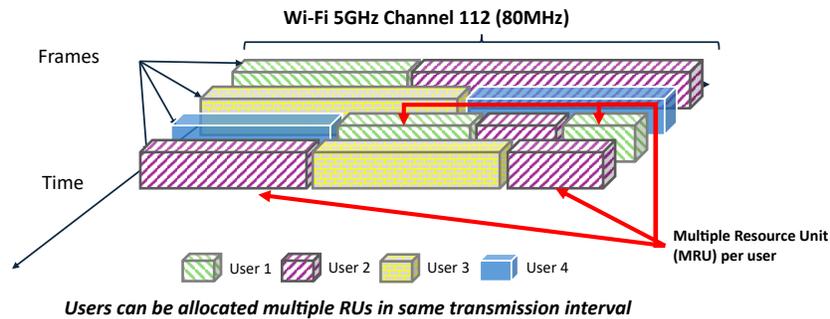


Figure 6, Multiple Resource Units (MRU)

Performance analysis demonstrates that MRU can reduce interference losses from 75% to 25%²⁴ compared to single RU allocation, effectively providing three times the data bandwidth availability in dense network environments. When multiple users transmit simultaneously with varying data length requirements, MRU can achieve 33% and 25% latency reduction compared to Wi-Fi 5 and Wi-Fi 6 respectively through optimized resource allocation algorithms.

4.6. Preamble Puncturing (Flexible Channel Utilization)

Preamble puncturing lets an access point disable specific 20 MHz subchannels within a wider channel (40/80/160/320 MHz) that are impacted by interference, while continuing transmissions over the remaining clean portions of spectrum. Previous Wi-Fi generations would basically make an entire channel unusable if there was interference present. This preserves much of the wide and ultra-wide channel throughput without forcing a full channel width to downgrade, for example 320 to 160 MHz, or 160 to 80 and works with OFDMA scheduling so punctured segments are excluded from resource unit assignments.

In 5 GHz DFS bands (UNII2/2e), puncturing cannot be used to bypass DFS rules – i.e. they still apply meaning if radar is detected on the primary 20 MHz channel, the device must vacate the channel. However, if radar is detected on a secondary 20 MHz segment within a bonded 80/160 MHz channel, puncturing may allow continued operation on the remaining nonradar segments, provided all DFS requirements are met - this can significantly improve overall performance by enabling that extra channel bandwidth

6 GHz does not use DFS for radar detection, so puncturing primarily addresses narrowband interferers and coexistence issues, especially valuable on 320 MHz channels where a single bad 20 MHz slice would otherwise degrade performance.

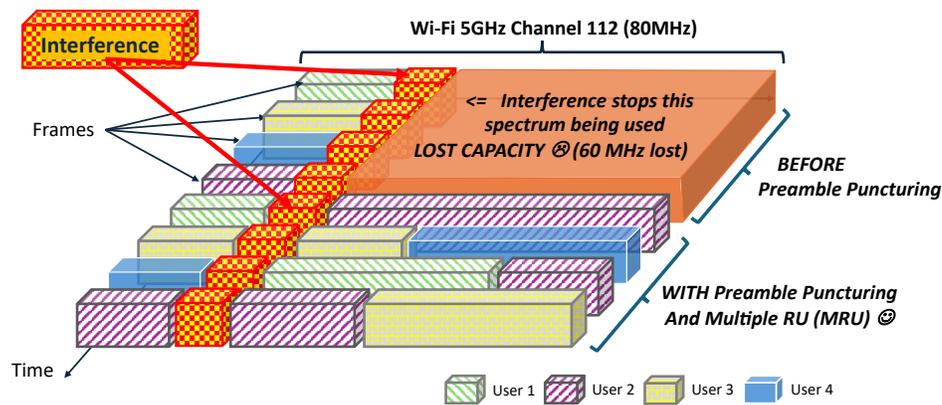


Figure 7, Multiple-RU and Preamble Puncturing

4.7. BSS Colouring and Spatial Reuse Groups (SRG)

While BSS colouring was added as a Wi-Fi 6 feature, and carried into Wi-Fi 7 without much change, it's worth identifying what it is, as Wi-Fi 8 makes major changes to it. At a high level this feature allows devices quickly tell “my BSS” versus “other BSS(OBSS)” by embedding a colour (1–63, basically an identifying number) into the Wi-Fi PHY header of transmitted frames. The feature allows concurrent transmissions by Wi-Fi devices active in the same channel in the same band. This enables an increase in throughput, reduced latency and better reliability.

While listening for Wi-Fi traffic, the device can determine if a frame is from its BSS or from an OBSS, simply by reading in the PHY header, matching the BSS colour, and either receiving the rest of the frame, or ignoring it. In the case of 50% OBSS traffic, this has a lot of benefits to a device from an energy saving perspective, allowing it to disregard a lot of frames.

In the case when the device wants to transmit to its own BSS, if it receives a frame during the Clear Channel Assessment (CCA), rather than giving up the chance to transmit, and reverting to the exponential backoff, it can decode the BSS colour of the PHY header and make a decision whether to transmit or not. If the BSS colour is from its own BSS, then the client will not transmit; alternatively it can detect the OBSS power level, and if below the configured OBSS-BD threshold, the device can go ahead and transmit. This does seem to contradict the CSMA/CA and CCA approach but it recognises the fact that

it's possible to have two transmissions on the same channel at the same time when two devices are talking to remote devices, and their transmissions will not corrupt each other.

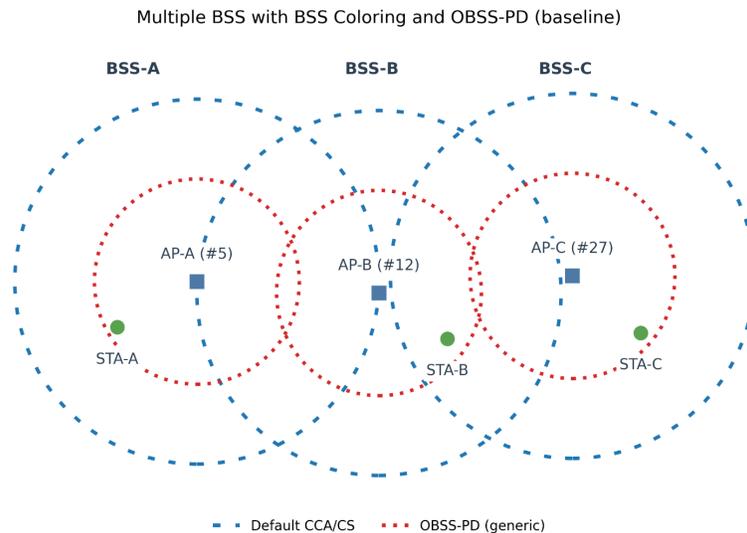


Figure 8, Basic BSS Colouring, Three APs, Three Different Colours

The feature also allows multiple related BSS to be added to a Spatial Reuse Groups (SRGs) - basically subsets/groups of neighbouring OBSSs for which a device/station (STA) uses a specific, possibly more aggressive SRG-OBSS-PD threshold and parameter. These inter-BSS frames are detected by BSS colour. There are different thresholds for SRG and non-SRG OBSS device, enabling more refinement in terms of concurrency decisions in dense deployments.

In practice, SR operation is the combination of BSS colouring for classification, OBSS-PD thresholding, and SRG/non-SRG policies to enable safe simultaneous transmissions when interference at intended receivers is unlikely to cause failures. The identification of OBSS into SRG groups is managed by external policy, such as a Wi-Fi management system.

Wi-Fi 7 integrated support for BSS-colouring and OBSS-PD/SRG into the updates for EHY PHY/MAC changes, but did not do anything significant.

4.7.1. BSS Colouring, SRG, and non-SRG - the cocktail party explanation

Imagine a large cocktail reception in one room.

- Each group of people chatting is like a Wi-Fi network (a BSS)
- Everyone shares the same room (same say 5 GHz channel), so groups can easily talk over each other if they don't coordinate

BSS colouring: coloured name badges for each group

- Give every conversation group a coloured badge: blue group, green group, and orange group
- If someone hears a voice from their own colour group, they pause so as to not talk over a teammate

- If they hear a different colour, they may decide it's just another group across the room and keep talking, as long as it sounds faint enough not to bother their own talking partner
- In Wi-Fi terms: the colour lets devices quickly tell “my network” vs “other network,” enabling informed decisions to be made when to Tx and when to wait

Spatial reuse with OBSS-PD: how loud the other group sounds

- Introduce a simple rule about “how loud is that other group?”
- If the other group is very loud nearby, be polite and wait
- If they're far away and just a murmur, it's fine to talk at the same time - maybe talk a bit more softly so you don't raise the room noise
- In Wi-Fi terms: devices use a configurable “detect” threshold for other group signal levels; if the other group is below that threshold, they are allowed to transmit concurrently, but with a restricted Tx power to stay polite

SRG vs non-SRG: friends vs strangers in the room

- Some of these groups are different than others. Some are “friendly neighbours” you coordinate with (SRG), others are just random strangers (non-SRG)
- For friendly neighbours (SRG), you agree on clearer sharing rules: if you hear them faintly, you can still talk, but you'll keep your voice lower than usual to avoid bothering them
- For strangers (non-SRG), you're more **conservative**: if you hear them at all, you're more likely to wait, or you'll speak only if they're really quiet and you still won't shout

In Wi-Fi terms: SRG uses a more aggressive threshold (easier to talk in parallel) but with stricter power limits; non-SRG uses a more conservative threshold with less aggressive reuse and looser power caps

Putting it together

- BSS colouring = the coloured badges that quickly tell “my group” vs “other groups.”
- OBSS-PD = the “how loud is the other group?” rule to decide if talking at the same time is okay, coupled with “lower your voice” when you do.
- SRG vs non-SRG = whether the other group is on your “friends list” (shared etiquette with more parallel talking but softer voices) or just a stranger (stricter politeness, fewer overlaps)

SRG vs non-SRG around My BSS (different OBSS-PD thresholds)

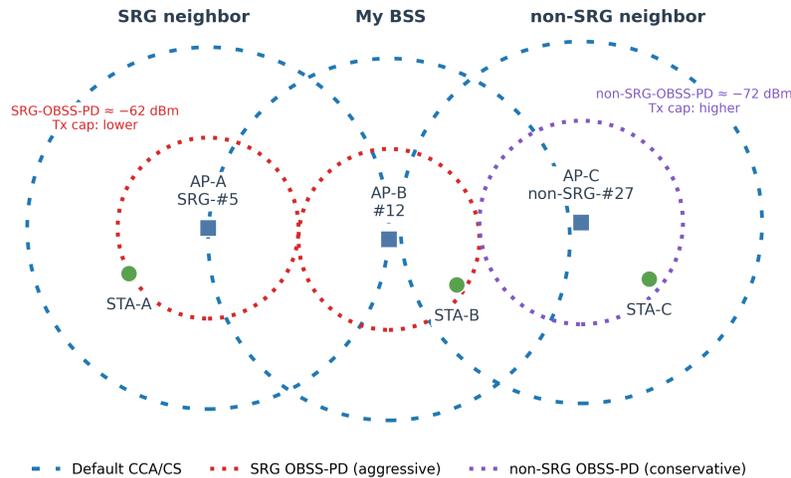


Figure 9, Complex BSS Colouring with SRG and Non-SRG groups

4.8. 802.11be feature mapping to Wi-Fi 7 Release 1

The following is an assessment of the mapping of 802.11be features to the final Wi-Fi 7 program

Feature	In IEEE 802.11be (EHT)	Wi-Fi CERTIFIED 7 Release 1	Notes
Multi-Link Operation (MLO)	Core MAC feature enabling simultaneous operation across multiple links/bands	Included as a headline capability to increase throughput, reduce latency, improve reliability	Program center piece; EMLSR/EMLMR explained in industry coverage
320 MHz channels (6 GHz)	EHT supports 320 MHz and related channelization in 6G Hz	Included where 6 GHz is available; advertised as doubling Wi-Fi 6's widest channel	Market enablement depends on local regulation and device/AP capabilities
4K-QAM (4096-QAM)	Defined modulation in EHT to raise PHY rates	Included as a headline feature for higher peak rates (~20% over 1024-QAM)	Wi-Fi program markets it as a headline boost
512 compressed Block Ack	EHT expands BA window to reduce MAC overhead	Included to improve efficiency and reduce overhead	Aligned optimization for bursty/high-throughput traffic
Multiple RUs to a single STA (MRU)	EHT defines MRU for OFDMA flexibility	Included to enhance spectrum efficiency and scheduling flexibility	Helps dense deployments via more flexible RU assignment
Triggered Uplink Access (UL)	EHT MU/TB uplink formats; UL scheduling enhancements	Included; optimization of Wi-Fi 6 triggered UL for latency-sensitive streams and QoS	Contributes to lower uplink latency and efficiency

Feature	In IEEE 802.11be (EHT)	Wi-Fi CERTIFIED 7 Release 1	Notes
Preamble puncturing / flexible channel use	Present in EHT to retain usable bandwidth under partial interference	Communicated under spectrum flexibility/operation in wide channels by vendors; part of EHT toolbox referenced alongside program features	Often bundled in vendor guidance with Wi-Fi 7 operation
Emergency Preparedness Communication Services (EPCS)	Within EHT MAC capability envelope	Included by the program to support NSEP services while maintaining priority/QoS	Public program item specifically called out in launch

Table 4, 802.11be Features in Wi-Fi 7 Release 1 and 2

4.8.1. Features in 802.11be not in Wi-Fi 7

The following features were discussed and are a part of the 802.11be definition but did not make it into the Wi-Fi CERTIFIED 7 program. A number of these will be expanded upon with 802.11bn and addressed in future Wi-Fi 8 specification.

Feature/Topic	802.11be definition	Notes versus WFA Wi-Fi 7
Expanded MUMIMO envelope and higher spatial stream ceilings (up to 16SS)	802.11be raises spatial stream limits and refines MU-MIMO beyond 802.11ax, enabling up to 16 spatial streams in the specification envelope	Widely documented as part of EHT; client devices generally support 1–2 streams, so maximum SS is typically platform dependent rather than a Wi-Fi program detail
Formal multi-AP coordination frameworks (beyond baseline)	Multi-AP coordination (e.g., coordinated/joint transmission) listed among 802.11be PAR candidate features	Candidate area in PAR discussions; more formal coordination frameworks are associated with subsequent standardization beyond 802.11be
Enhanced link adaptation/retransmission (HARQ)	HARQ identified as a candidate enhancement for robustness and reliability in 802.11be PAR notes	Appears in candidate feature lists rather than certification headlines; emphasis in Wi-Fi 7 focuses on MLO/320 MHz/320MHz/MRU/4KQAM
Forward compatibility constructs (e.g., frame formats/USIG/PHY details)	802.11be introduces frame formats with improved forward compatibility and related PHY signalling constructs	Important to the standard’s evolution but typically not marketed as a separate Wi-Fi 7 certification feature
Deep TSN alignment (802.1Q deterministic scheduling integration)	PAR candidate items include integration with TSN extensions (802.1AS timing, bounded latency shapers, etc.) to support low latency real-time traffic	Positioned as candidate/adjacent areas in PAR text; not part of Wi-Fi 7 program marketing
Detailed channelization variants beyond marquee 320 MHz/320MHz (e.g., 240 MHz/240MHz, mixed contiguous/non-contiguous)	802.11be supports additional BW/channelization options (320/240 MHz/240MHz and non-contiguous combinations) beyond the headline 320 MHz/320MHz	The standard’s channelization flexibility exceeds the Wi-Fi 7 program

Table 5, 802.11be Features not included in Wi-Fi CERTIFIED 7

5. 802.11bn/Wi-Fi 8 Features

5.1. Overview

IEEE 802.11bn Ultra High Reliability brings a change in Wi-Fi connectivity. Previous versions of Wi-Fi exploited speed improvements through a variety of features like 160/320 MHz channels, 1024 and 4096-QAM, Multi Link Operation (MLO), more spatial streams, etc., resulting in the ability to deliver 46 Gbps (this is more likely to be ~19Gbps for residential APs). IEEE 802.11bn work is focussed on ultra-high reliability, with a view to deliver deterministic performance, reductions in latency, jitter and packet loss in dense and demanding environments. Large aspects of 802.11bn will be carried forward into the Wi-Fi Alliance Wi-Fi 8 certification program, providing the basis for the deployment of interoperable Wi-Fi 8 devices. For residential service providers, Wi-Fi 8 represents more of a strategic enhancement rather than an overhaul of existing ecosystems dominated by Wi-Fi 5, 6, and 7 devices.

Wi-Fi 8 addresses the “reliability gap” - that disconnect between the theoretical speeds and real world performance that exists in mixed device networks. With about 20+ billion Wi-Fi devices in existence, and about 7% using Wi-Fi 7 as of 2024²⁵, legacy populations remain the norm. Service providers face challenges like Multi-Dwelling Units (MDU) interference, coverage dead zones in Single Family Units (SFU), and unpredictable latency for applications. Wi-Fi 8 features intend on bridging this gap through coordination and QoS mechanism, working with existing deployments, while unlocking significant benefits through Wi-Fi 8.

Wi-Fi 8 is designed to be fully backward compatibility. Like previous Wi-Fi generations it is expected to enhance rather than replace existing infrastructure based on Wi-Fi 5/6/7. Operator experience with some Wi-Fi 7 features has not gone as smoothly as hoped with some interop issues with MLO and operational issues with WPA3, especially given legacy client mixes. The desire by operators is to ensure Wi-Fi 8 certification will ensure clean interoperability from the start, and a smoother deployment path from day #1, while addressing those stubborn Wi-Fi 7 issues.

Given the intense competition in the residential service provider space for customer acquisition and retention, as Wi-Fi 8 enters the final phases of specification and certification, the service providers will begin marketing campaigns around the superior Wi-Fi delivered by Wi-Fi 8.

Initial focus will be on upgrading existing Wi-Fi 5/6/7 single AP homes with a new Wi-Fi 8 AP. The new AP will of course work with all legacy Wi-Fi devices in the home, along with any existing Mesh APs that may already be in place (assuming Wi-Fi management is properly implemented/etc.). Legacy devices will experience the benefit of some Wi-Fi 8 features and hardware capability immediately. In most cases, Wi-Fi 8 adoption by subscribers will begin on big ticket devices, like new phones, tablets, laptops, etc. connecting to the new Wi-Fi 8 APs. These new devices will either replace existing devices or those devices will trickle down to other users in the home, with some much older devices being removed. These new Wi-Fi 8 clients will immediately enable some of the more advanced features of Wi-Fi 8. The side benefit of new Wi-Fi 8 devices being far more efficient and reliable is network capacity is freed up that helps the performance of all the older devices.

Obviously some homes will receive a Wi-Fi 8 mesh network, adding extenders to a new gateway or main AP. This enables all the cool features in Wi-Fi 8 based around Multi-AP Coordination (MAPC), leading to a major increase in network efficiency and performance both for the mesh and Wi-Fi 8 devices, and even indirectly for legacy devices.

As more Wi-Fi 8 clients connect to the network, the operation and reliability of the network will continue to improve, compared to when it was based on Wi-Fi 7 or earlier generations.

The following sections dig deep into each of the key features of 802.11bn, which is the precursor to the Wi-Fi Alliance Wi-Fi 8 release. That WFA release will consist of a subset of the features 802.11bn offers, determined by the WFA members. These features are developed around baseline capabilities and more sophisticated options; the WFA members will determine the scope of each of the selected features that can be incorporated into the available timelines to hit the end of 2027 certification launch date. The selection of features and capabilities of those features is not known at this stage, but a lot of Wi-Fi silicon providers are discussing many of these features, and expressing interest to have them incorporated. IEEE 802.11bn features will be introduced and picked apart below, with an explanation of how they work with legacy devices and Wi-Fi 8 devices, and what their benefits are. Coordination based features deliver significant performance improvements with Wi-Fi 8. In most cases this requires a central coordination function to be active on one of the APs is coordination, and that necessitates a level of local control between the different APs in a Wi-Fi network to manage the required coordination. In most scenarios an AP will be selected to perform the required coordination function, driving the operation of the network. Most coordination features need a high quality backhaul link between participating APs, with some of the features requiring extremely tight timing.

5.2. Co-ordinated Spatial Reuse (Co-SR)

Spatial Reuse was outlined already, and is an effective means of improving operation of Wi-Fi in the cases where there are overlapping BSS that are operating on the same channel within a band. The use of BSS colouring to identify the different BSSs and the different OBSS-PD thresholds for power detection are key to the feature, as is the need for distinguishing between BSS in a SRG (defined and communicated to all members in the SRG) and those BSS not in the BSS. The solution in Wi-Fi 6 and 7 is pretty basic as defined and brought some efficiencies but as it operating on the basis of independent APs making their own decisions without any sharing of information with other APs, there was obvious improvements possible.

Co-ordinated Spatial Reuse (Co-SR) in 802.11bn/Wi-Fi 8 transforms spatial reuse from a reactive, independent AP feature into a proactive, intelligent, multi-AP coordinated capability. Co-SR works with existing Wi-Fi clients and Wi-Fi 8 clients, meaning the benefits of SR are amplified resulting in much more substantial throughput gains, much lower latency and far better reliability.

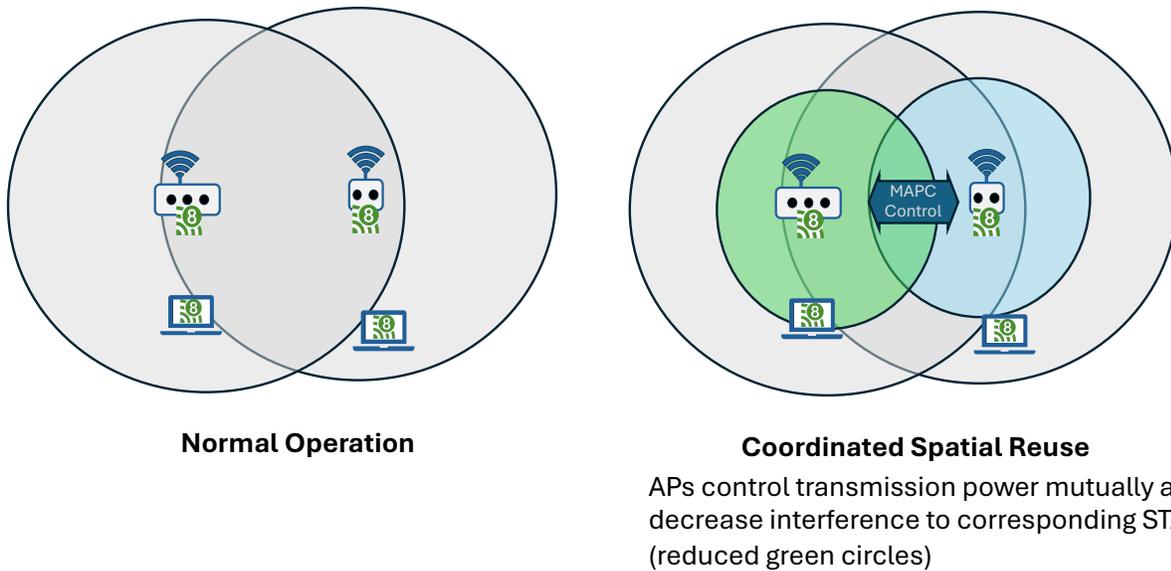


Figure 10, Co-ordinated Spatial Reuse (Co-SR)

The feature itself depends on multi-AP signalling, synchronised power control and AP interconnection (either wired or wireless backhaul). It works with (but not completely dependent on) the other Wi-Fi 8 features such as Co-Time Division Multiple Access (TDMA), Co-Beamforming, and Distributed-Multi-Link Operation (MLO) to deliver maximum benefit. Legacy devices cannot trigger Co-SR, but Wi-Fi 8 clients can receive and act upon Co-SR trigger frames sent by APs to adjust transmission parameters and transmit simultaneously (which is the key to the feature).

The key elements of Wi-Fi 8 Co-SR are

- **Multi-AP Coordination Framework (MAPC):** APs exchange interference and buffer status via MAPC-ICF/ICR control frames.
- **Synchronized Power Control:** All APs in a coordination group agree on reduced TX power levels to enable concurrent transmissions without harmful interference.
- **Scheduled Concurrent TXOPs:** C-SR aligns TXOP start times across APs, ensuring parallel transmissions at designated intervals.
- **Real-Time Interference Sharing:** APs periodically share fine-grained interference measurements to adapt reuse decisions dynamically.
- **AI-Driven Optimization (Optional):** Reinforcement-learning algorithms choose optimal AP–STA pairs for spatial reuse under dynamic conditions.

5.3. Co-ordinated Beam Forming (Co-BF)

Traditional beamforming in Wi-Fi 6 and 7 operates on a single basic service set (BSS), where an AP with multiple antennas can focus its own RF energy toward a specific client device. This spatial focusing improves the signal-to-noise ratio (SNR) at the target receiver and reduces spillover energy in other directions. However, single-AP beamforming has limitations: it cannot account for interference to devices served by neighbouring APs, and it provides no mechanism for coordinated interference cancellation across multiple APs sharing the same channel.

Co-BF works by extending multi-user MIMO (MU-MIMO) concepts across AP boundaries. In single-AP MU-MIMO, an AP simultaneously serves multiple clients by creating spatially separated beams, with nulls formed in directions where they would cause interference. Co-BF applies this principle at the network level: multiple APs work together to create constructive beams toward their intended clients while forming nulls toward clients served by other APs in the coordination group.

The Co-BF process begins with cross-BSS channel sounding, where each AP and its associated stations perform channel measurements not only within their own BSS but also with stations and APs from neighbouring BSSs. This channel state information (CSI) collection reveals the spatial channel details between every AP-station pair in the coordination group. The CSI data includes both amplitude and phase information for each spatial stream, providing the raw data for joint beamforming calculations.

CSI exchange over the backhaul network follows, where APs share their collected channel measurements through high-speed wired or wireless connections. This exchange must happen quickly enough to remain relevant - the channel conditions change rapidly in mobile environments - requiring low-latency backhaul infrastructure. The shared CSI dataset enables each AP to understand not just its own channel conditions but the interference relationships throughout the entire coordination group.

Joint steering matrix computation represents the mathematical heart of Co-BF. Using the combined CSI from all APs and stations, complex algorithms calculate per-AP precoding matrices that optimize network-wide performance. Each AP's steering matrix is designed to maximize signal power at its intended receivers while minimizing interference at receivers served by other APs. This joint optimization often requires significant computational resources and must balance conflicting objectives across multiple links. This might open the door to some AI/ML handling in the future.

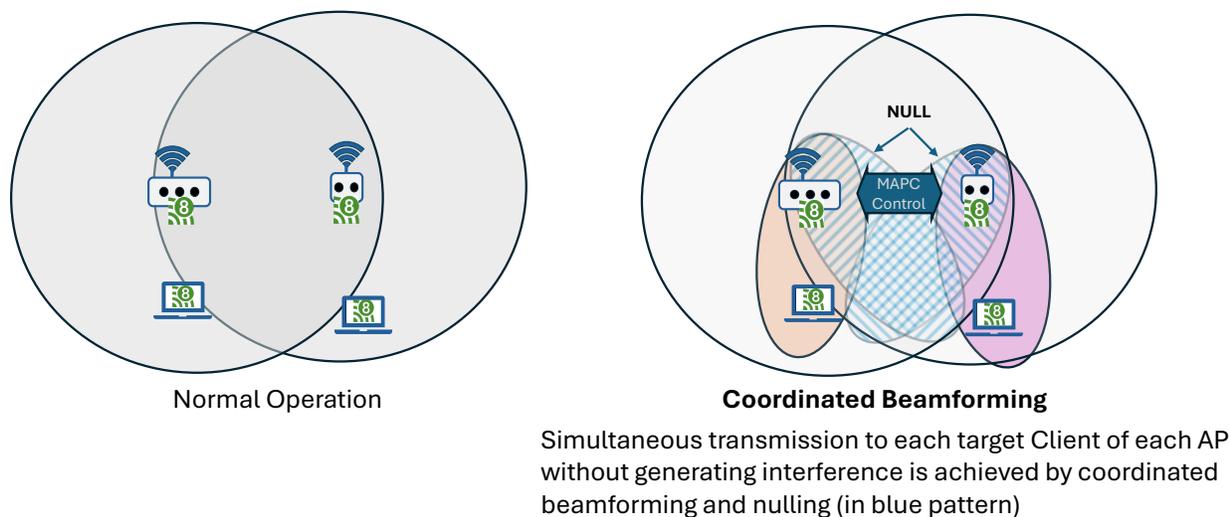


Figure 11, Co-ordinated Beam Forming (Co-BF)

Synchronized transmission execution demands precise timing coordination, typically requiring sub-microsecond alignment between participating APs. When APs transmit their pre-coded signals simultaneously, the overlapping RF patterns create constructive interference at target receivers and destructive interference (nulls) at protected receivers. This synchronization is challenging to maintain, particularly over wireless backhaul connections where timing jitter can disrupt the carefully orchestrated interference patterns.

The requirements for ultrafast matrix computation and distribution of precoding matrices to each AP requires a lot of fast backhaul. This maybe challenging in the case of residential homes dependent on wireless backhaul connectivity, but maybe offset by the relatively static configuration of a residential home were 70-80% of devices are statically positioned.

Maximum Co-BF performance gains are achieved in environments with major spatial separation between APs and their clients. Field measurements show 3–8 dB SINR improvements¹¹ at stations where interference from neighbouring APs traditionally limits performance. In dense mesh networks, aggregate throughput improvements of 20–50 % are common as Co-BF enables more aggressive frequency reuse than traditional beamforming allows.

The implementation complexity of Co-BF represents both a challenge and an opportunity. While the CSI exchange and joint computation require sophisticated hardware and software, the resulting performance improvements justify the investment in dense enterprise and high-end residential deployments. Co-BF transforms beamforming from a local optimization technique into a network-wide interference management fabric, marking a significant evolution in Wi-Fi's spatial processing capabilities. The limitations of Wi-Fi backhaul, and matrix computation overhead, particularly in residential settings might impact on Co-BF deployments

5.4. Co-ordinated Time Division Multiple Access (Co-TDMA)

In Wi-Fi 6/7, contention over the medium - even with OFDMA and MU-MIMO - remains unpredictable, resulting in varying backoff delays and collision rates. Coordinated Time Division Multiple Access (Co-TDMA) in Wi-Fi 8 operates by reserving a shared TxOP and slicing it into fixed time slots assigned to each AP in a coordination group, providing a level of determinism.

Upon initiating a Co-TDMA session, APs agree on the total TxOP duration and slot lengths. A MAPC control frame carries the slot assignment map, which each AP uses to configure its transmission schedule. During its allocated slot, an AP transmits uninterrupted; during other slots it remains silent, foregoing contention entirely. Legacy CSMA/CA continues to be used by any non-scheduled, best-effort traffic outside the reserved TxOP window.

By removing inter-AP collisions, Co-TDMA can reduce tail-latency by up to 95%¹⁸ and improve channel efficiency by 20-30% in coordinated environments²¹. This deterministic access model is particularly advantageous for real-time control, and AR/VR applications that demand predictable timing and minimal jitter.

The carrier sense multiple access with collision avoidance (CSMA/CA) mechanism that has governed Wi-Fi channel access since 802.11's inception becomes increasingly problematic as network density grows. Under CSMA/CA, each AP must independently win access to the channel through a random backoff process, competing with all other APs within range. While this distributed approach scales reasonably well with moderate AP density, it suffers from the hidden node problem, exponentially increasing collision probability, and unpredictable worst-case delays that can exceed hundreds of milliseconds. This also is amplified by the operation of mesh networks sharing common backhaul channels, increasing the likelihood of delay, as each mesh hop duplicated the bandwidth demand from each remote fronthaul.

In Wi-Fi 6 and 7, OFDMA provides some relief by allowing an AP to serve multiple clients simultaneously within its own BSS, but inter-AP coordination remains governed by CSMA/CA. Multiple

APs sharing the same channel still must contend for transmission opportunities, often resulting in one AP monopolizing the channel while others defer. This access pattern wastes significant channel capacity and creates unpredictable latency patterns that frustrate time-sensitive applications.

Co-TDMA in Wi-Fi 8 fundamentally changes this paradigm by replacing probabilistic access with deterministic scheduling. The concept draws inspiration from traditional cellular time-division systems but adapts to Wi-Fi's distributed architecture and backward compatibility requirements. Rather than requiring a centralized base station controller, Co-TDMA operates through distributed coordination among participating APs, with one AP assuming temporary coordinator responsibilities for each shared TXOP.

The coordination process begins when APs negotiate to form a Co-TDMA group and agree on a shared TXOP duration. The TXOP is then subdivided into discrete time slots, with each participating AP assigned specific slots based on its traffic requirements and QoS needs. A coordination frame carries these slot assignments to all group members, establishing a collision-free transmission schedule for the duration of the TXOP. This deterministic allocation eliminates the random delays and collisions inherent in CSMA/CA.

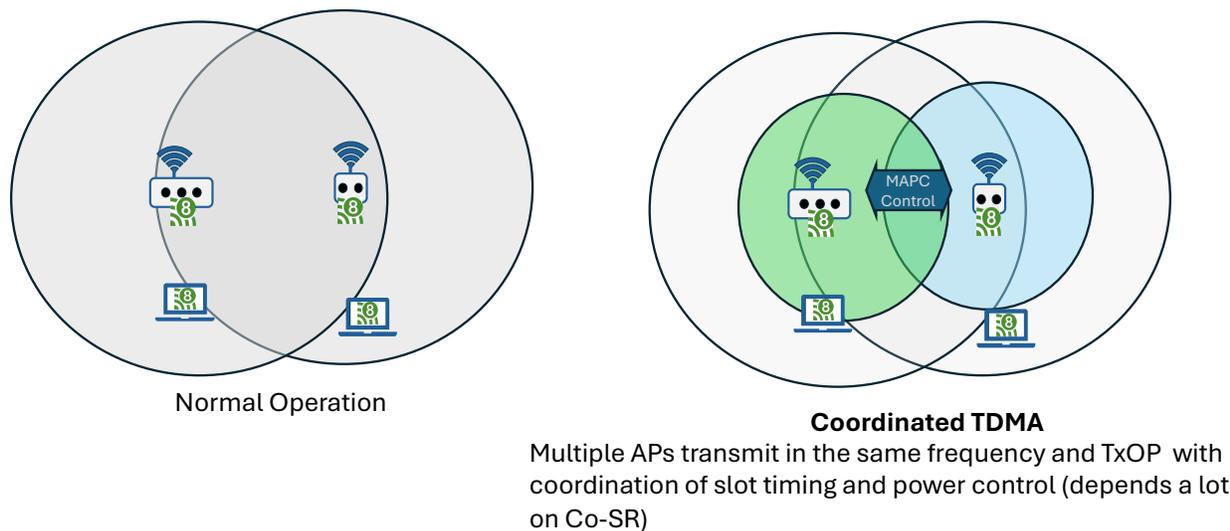


Figure 12, Co-ordinated TDMA

Slot timing synchronization represents a critical implementation challenge, as participating APs must maintain precise time alignment to prevent slot boundary violations. Synchronization accuracy requirements typically demand sub-microsecond precision, achievable through high-quality reference clocks or precision time protocol (PTP/IEEE1588) over wired backhaul. Wireless backhaul connections introduce additional timing jitter that must be compensated through adaptive synchronization algorithms.

The hybrid access mechanism preserves backward compatibility by allowing CSMA/CA operation outside the coordinated TxOPs. Legacy stations and APs that don't support Co-TDMA continue to operate normally during uncoordinated periods, while Co-TDMA-capable devices gain the benefits of collision-free access during scheduled intervals. This coexistence approach enables gradual deployment without requiring immediate network-wide upgrades.

Co-TDMA's deterministic nature particularly benefits latency-sensitive applications that suffer under CSMA/CA's unpredictable access delays. Industrial automation protocols, real-time gaming, AR/VR

applications, and voice communications all benefit from the bounded, predictable latency that Co-TDMA provides. Measurements show worst-case latency reductions of up to 95% compared to CSMA/CA operation under high load conditions.

Channel efficiency improvements stem from the elimination of collision-related overhead and the reduction of inter-frame spacing requirements within coordinated TxOPs. Without the need for random backoff periods and collision recovery mechanisms, more airtime becomes available for actual data transmission. Typical efficiency gains range from 20% to 30% in coordinated deployments, with larger improvements possible as the number of coordinated APs increases. Co-TDMA represents a significant step toward the deterministic, reliable communication that defines Wi-Fi 8's ultra-high reliability vision.

5.5. Coordinated Restricted Target Wake Time (Co-RTWT)

Coordinated Restricted Target Wake Time (Co-RTWT) in 802.11bn optimises power management and interference mitigation across multiple APs operating in overlapping basic service sets (OBSS). Co-RTWT enables APs to coordinate their restricted Target Wake Time (rTWT) schedules or protect each other's rTWT schedules, ensuring time-aligned transmission opportunities to reduce contention and improve network efficiency in dense residential or enterprise deployments.

At its core, Co-RTWT involves one AP negotiating with another to extend protection for its rTWT schedules, through management frames exchange scheduling parameters. Once negotiated, the AP that provides protection guarantees that its TXOP ends before the start of the rTWT service periods (SPs) of the other AP. This approach prevents overlapping transmissions that would cause interference.

APs also the coordinated OBSS rTWT schedule in their beacon frames to inform all connected clients of the baseline rTWT rules. This helps clients remain synchronised. This coordination is essential for dense environments with multiple overlapping APs to ensure smooth traffic flow, reduce collisions, and extend device battery life, particularly for IoT and low-power client devices.

Overall, Co-RTWT in 802.11bn represents a significant improvement in multi-AP cooperation for Wi-Fi networks, focussing on dynamic schedule negotiation and protection to enhance coexistence and reliability. It allows Wi-Fi 8 to deliver ultra-high reliability and power-efficient operation in complex AP networks, minimising interference and supporting deployments of large numbers of client devices.

5.6. Distributed Multi-Link Operation (dMLO) - Seamless Mobility

Traditional Wi-Fi roaming in legacy networks requires a complete disassociation and reassociation when clients move between APs. This process is followed by renegotiation of security parameters and quality-of-service settings. Even with 802.11r fast roaming optimizations, this handoff process typically requires 50-200 milliseconds, causing noticeable interruptions in voice calls, video streaming, and real-time applications.

Wi-Fi 6 and 7 introduced Multi-Link Operation (MLO) within single physical devices, allowing clients to simultaneously connect to multiple bands of the same AP. While this single-AP MLO improves performance and provides resilience against band-specific interference, it doesn't address the fundamental roaming problem when clients move beyond the coverage area of their current AP. The association and security context remain tied to a specific physical device, requiring the traditional roaming process when moving between APs.

Distributed MLO (dMLO) in Wi-Fi 8 extends the MLO concept across multiple physical APs by presenting them as a single logical Multi-Link Device (MLD). Each participating AP advertises the same

Multi-Link Identifier (MLID) in its beacon frames, signalling to clients that these APs function as components of a unified device. From the client's perspective, associating with any AP in the dMLO group establishes connectivity to the entire logical device, enabling seamless movement between physical APs without re-association. Note that the CableLabs Mobile-Wi Fi feature provided similar functionality to how dMLO handles steering and achieved sub 5msec steering times.

A state synchronization mechanism lies at the heart of dMLO operation. All APs sharing an MLID must maintain synchronized copies of client association state, including security keys, QoS parameters, power management settings, and scheduled wake times (TWT). This synchronization occurs over the backhaul network connecting the APs, requiring reliable, low-latency communication to ensure state consistency across all group members. Changes to client state must be rapidly propagated to prevent inconsistencies that could disrupt service.

Link steering algorithms dynamically route client traffic over the optimal AP link based on real-time performance measurements. Unlike traditional roaming decisions based on received signal strength alone, dMLO can consider multiple factors including channel utilization, interference levels, client battery status, and application QoS requirements. The network can proactively shift traffic before link quality degrades, maintaining optimal performance without client intervention.

Traffic continuity during handoffs represents a key dMLO advantage. Because all APs in the group maintain the client's security and QoS context, traffic can be seamlessly redirected from one AP to another without interrupting ongoing sessions. Buffered frames can be forwarded over the backhaul to ensure no packet loss during the transition. This continuity is particularly valuable for voice-over-IP calls, video conferences, and other applications sensitive to brief interruptions.

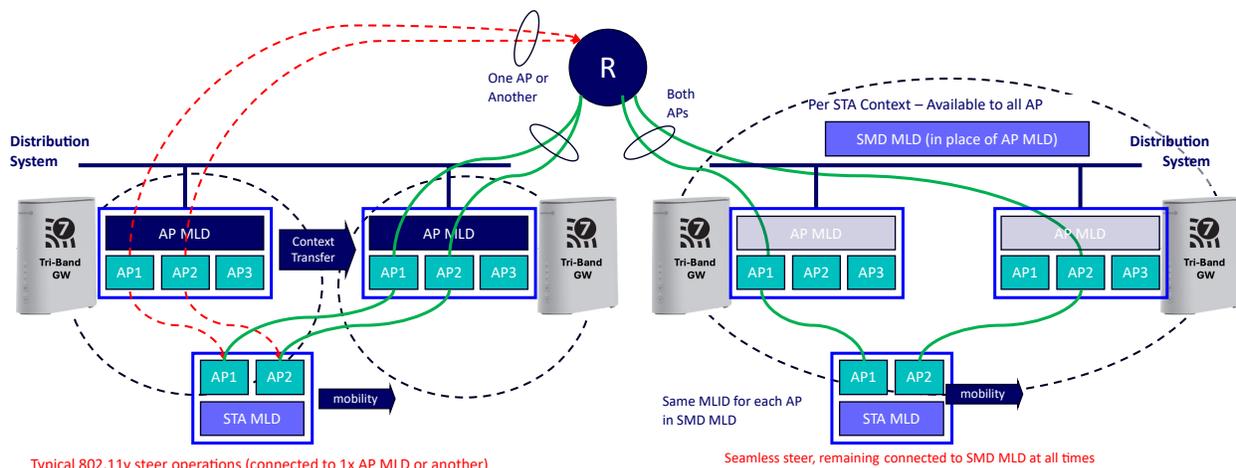


Figure 13, MLO and distributed-MLO Operation (Seamless Mobility Domain)

The implementation complexity of dMLO requires sophisticated coordination protocols and state management mechanisms. APs must implement distributed databases to maintain synchronized client state, conflict resolution algorithms to handle simultaneous updates, and failure recovery procedures to maintain service when individual APs become unavailable. The backhaul network becomes a critical component, as its reliability and latency directly impact dMLO performance.

Performance measurements demonstrate dMLO's significant impact on mobile user experience. Handoff interruptions drop from tens of milliseconds to sub-millisecond levels, while session continuity exceeds 99% for properly configured deployments. Voice call quality remains high during movement, video streaming continues without rebuffering, and real-time applications maintain their interactivity. dMLO

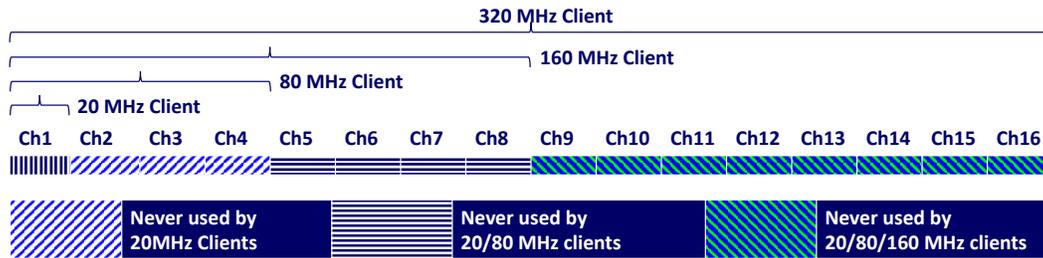


Figure 14, DSO sub-channel options

Studies have shown between 20 and 80% throughput gains¹⁶ under heavy load for Wi-Fi 8 devices, unfortunately Wi-Fi 6 and 7 devices cannot avail of DSO. However these legacy clients can benefit from Wi-Fi 8 devices using DSO, resulting in fewer collisions and lower latency of client devices sharing the primary channel.

5.7.2. Non-Primary Channel Access (NPCA)

Wi-Fi communications between APs and client devices always begins with gaining access to the Primary 20 MHz sub-channel of a 20/40/80/160/320 MHz channel before transmitting data. If this channel is busy due to narrow-band interference (OBSS transmissions, etc.), then the usual CSMA/CA process forces the devices to backoff and retry their access. Non-Primary Channel Access (NPCA) enables an AP and client to opportunistically transmit over secondary channels when their primary 20 MHz sub-channel is congested, with a throughput boost of up to 37.5% realistic environments¹⁰ and delivering more consistent low-latency performance.

In a dense or mixed use environment, overlapping APs (OBSS) can congest a shared primary channel, resulting in the adjacent spectrum being underutilized - think of three APs with 40/80/160 MHz channels but the 20 MHz primary channel is shared/overlapping. This forces these APs to share the airtime equally, degrading effective throughput for all users. For example an AP with 160 MHz channel and an 80 MHz client is limited to the narrower channel, effectively wasting 50% of its capacity.

NPCA can address this by allowing devices to switch to alternative sub-channel with the wider bandwidth if the primary sub-channel is occupied by interference/OBSS. In the example just mentioned this would allow the 80 MHz client device to use the sub-channel at 80 MHz as if it were its primary sub-channel at the start of the 160 MHz AP bandwidth.

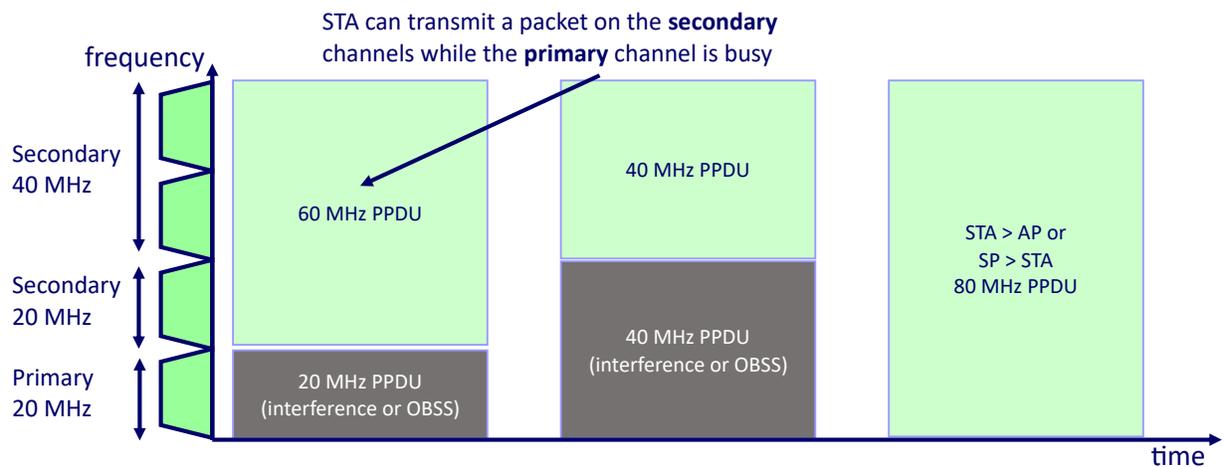


Figure 15, NPCA Channel Switching Options

NPCA works by clients and APs exchanging capabilities during (re)association/channel update time, identifying the supported sub-channels and the channel switching latencies. The AP continuously monitors the primary sub-channel for Co-Channel Interference (CCI) from other APs and if interference levels exceed a preconfigured threshold it triggers NPCA. The AP and client jointly select an idle “NPCA primary channel” (another 20 MHz channel within the 160/320 MHz channel bandwidth). Within the ongoing TxOP, the AP instructs the client to switch/tune to the NPCA channel, where they complete the PPDU on the alternate primary without releasing the TxOP. Once complete, they return to the original primary channel for normal operation.

5.7.3. Dynamic Bandwidth Expansion (DBE)

In legacy Wi-Fi 6 and 7 networks, an access point (AP) and its clients must agree on a single channel bandwidth for the duration of their association or a target wake time (TWT) agreement. Once that bandwidth-20, 40, 80, 160, or 320 MHz-is established, the AP and client remain confined to those subcarriers, even if the client’s capabilities or RF conditions change. In practice, many clients (especially battery- or cost-constrained devices) support only narrower bandwidths-typically up to 80 or 160 MHz-while the AP can operate on a much wider 320 MHz channel. This mismatch forces the AP to down-scale its entire physical-layer data unit (PPDU) to the client’s narrower bandwidth, wasting potentially hundreds of megahertz of spectrum and limiting aggregate throughput.

Dynamic Bandwidth Extension (DBE) in Wi-Fi 8 resolves this static-bandwidth constraint by allowing an AP to dynamically reconfigure the active channel bandwidth on a per-TXOP basis to match each client’s instantaneous capability and demand-without renegotiating the underlying association or TWT agreement. When preparing a downlink transmission to a narrow-bandwidth client, the AP confines the PPDU to the client’s supported sub-set of the full 320 MHz channel. Immediately afterward, the AP can revert to the full 320 MHz for subsequent transmissions to wider-capable clients, all within the same TXOP if desired.

DBE operates in two phases. During the TxOP’s initial control phase, the AP queries each client for its current bandwidth capability and switching latency (the time required to retune its radio from one sub-band to another). The AP then composes a composite PPDU schedule that allocates specific sub-band widths to each client’s data unit, ordered to minimize switching overhead. In the DBE transmission phase, the AP and clients swiftly retune their radios to the assigned sub-bands, exchange data, and retune again for the next client or the full-band transmission. Client radios employ fast-tuning hardware and optimized

DSP pipelines to achieve sub-100 μ s retune times, ensuring that the overall TXOP duration remains efficient.

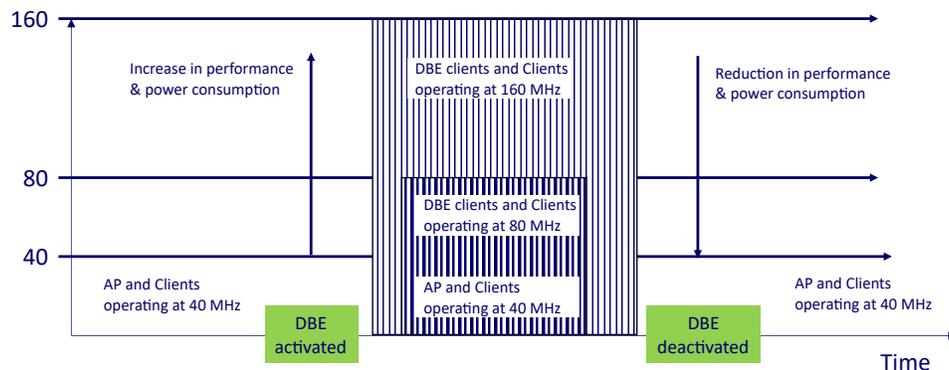


Figure 16, Dynamic Bandwidth Expansion (DBE)

By tailoring each PPDU's bandwidth slice to the client's capabilities and the instantaneous traffic pattern, DBE delivers multiple benefits:

- Spectral Efficiency - The AP uses the full 320 MHz channel for high-capability clients, while narrow-band clients consume only the spectrum they need, maximizing total network capacity.
- Latency Reduction - Clients waiting for narrow-band transmissions no longer delay full-band PPDUs to other clients; each link transmits at its optimal width.
- Energy Savings - Narrow-band clients spend less time transmitting and receive only within their tuned sub-bands, reducing radio-on time.

Simulation and prototype studies show that DBE can boost aggregate downlink throughput by 10% to 30% in mixed-capability networks and reduce per-client latency variance by up to 25% compared to static-bandwidth operation. Crucially, DBE requires only AP firmware enhancements and modest client-side DSP support for fast retuning-making it a low-barrier, high-impact feature for both enterprise and advanced residential deployments.

5.7.4. Distributed Resource Units (DRU)

Traditional OFDMA in Wi-Fi 6 and 7 allocates contiguous blocks of subcarriers (Resource Units or RUs) to each client device, with the client's transmit power distributed evenly across all assigned tones. While this approach works well for high-SNR scenarios, it becomes problematic when clients operate at the cell edge or in interference-limited conditions. Small RUs like RU26 (26 tones) provide fine scheduling granularity but suffer from low per-tone power spectral density, limiting their effective range and reliability.

The power spectral density limitation stems from regulatory constraints on total transmit power. Whether a client transmits on 26 tones or 996 tones, its total power remains constant (typically 20-23 dBm). When spread across a large RU, each tone receives only a fraction of the total power. For edge devices or IoT sensors with limited power budgets, this power spreading can push per-tone SNR below the threshold required for reliable communication, forcing the use of very conservative modulation and coding schemes.

DRU in Wi-Fi 8 breaks the contiguous allocation constraint by allowing clients to concentrate their full transmit power on non-contiguous, distributed tones across the available bandwidth. Instead of spreading 20 dBm across 996 tones (yielding approximately 20 μ W per tone), a client can concentrate the same 20 dBm on just 26 distributed tones, achieving nearly 800 μ W per tone—a 15.8 dB power spectral density improvement. This power concentration dramatically extends uplink range and reliability for challenging link conditions.

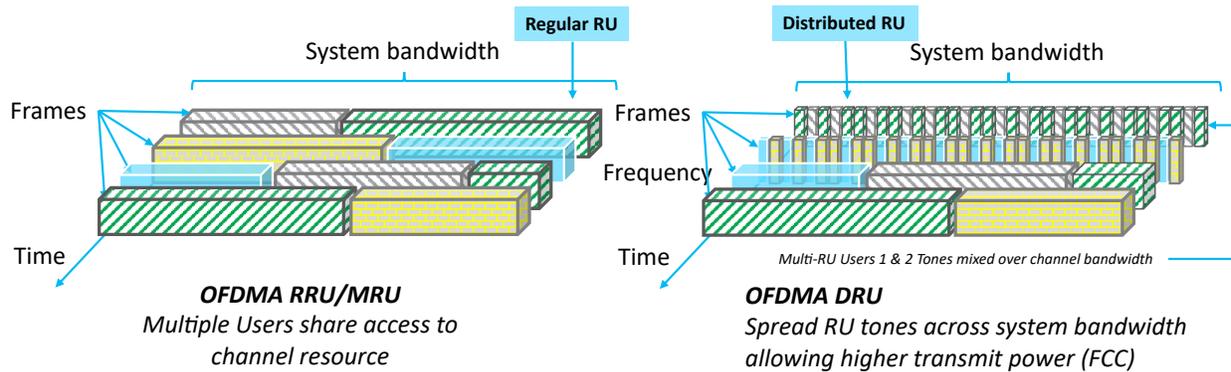


Figure 17, RRU/MRU versus Distributed RU (DRU)

The DRU allocation process begins during client association, when devices report their supported RU sizes and power capabilities through management frames. The AP maintains a database of each client's DRU capabilities and current link conditions. When uplink traffic requires enhanced range or reliability, the AP can issue a DRU grant specifying exactly which non-contiguous tones the client should use for its transmission. The grant includes both the tone indices and any power control parameters needed for interference management.

RU Size	RRU	DRU (M below indicates the number of tones in each 13 subcarriers (1 MHz))			
		DBW = 20 MHz	DBW = 40 MHz	DBW = 80 MHz	DBW = 160 MHz
26	2.08	10.14 (M=2)	13.15 (M=1)	13.15 (M=1)	13.15 (M=1)
52	5.09	11.39 (M=3)	13.15 (M=2)	16.16 (M=1)	16.16 (M=1)
106	8.18	11.47 (M=6)	14.48 (M=3)	16.24 (M=2)	19.25 (M=1)
242	11.77	N/A	14.39 (M=7)	16.82 (M=4)	19.83 (M=2)
484	14.78	N/A	N/A	17.40 (M=7)	19.83 (M=4)
996	17.91	N/A	N/A	N/A	20.53 (M=7)

Figure 18, Maximum Transmit Power of RU (dBm)

Client-side DRU implementation requires sophisticated PHY processing to handle non-contiguous subcarrier mapping. Traditional OFDMA processors expect contiguous tone allocation, but DRU demands the ability to map data streams onto arbitrary subcarrier patterns. This processing complexity increases with the degree of tone distribution and the total bandwidth span, requiring more advanced digital signal processing capabilities than conventional OFDMA operation.

The performance benefits of DRU are most pronounced for small RU allocations and challenging link conditions. RU26 allocation with DRU can extend uplink range by 3–5 dB¹⁹ compared to contiguous RU26, often making the difference between reliable communication and frequent retransmissions. IoT devices, outdoor sensors, and basement-located smart home devices particularly benefit from DRU's

range extension capabilities. Retransmission rates typically drop by 20% to 40% in low-SNR environments.

Interference management with DRU requires careful coordination to prevent distributed transmissions from interfering with neighbouring cells. The AP must consider not only the local interference environment but also the potential impact of distributed tones on adjacent networks. This coordination becomes particularly important in dense deployment scenarios where multiple APs might simultaneously use DRU for their respective clients.

Legacy device coexistence represents both a challenge and an opportunity for DRU deployment. Wi-Fi 6 and 7 clients cannot use DRU themselves, but they benefit indirectly when Wi-Fi 8 clients use DRU for their transmissions. By concentrating power on distributed tones rather than spreading it across full-band allocations, DRU reduces the overall interference footprint, leaving more clean spectrum available for legacy devices on the primary channel.

5.7.5. New MCS

802.11bn identified the benefit of including new finer grained Modulation and Coding Scheme (MCS) options to address a number of issues identified with the existing range of MCS values in 802.11be/Wi-Fi 7. The new MCS only focused on the “coding” aspects of MCS, and avoided changing any modulation options. The key benefit was to ensure a smoother transition between MCS data rates, by closing the large sensitivity Signal to Noise Ratio (SNR) gap between some existing adjacent MCSs (3.5 dB between MCS2 and MCS3, and 4+ dB between MCS4 and MCS5).

Mod/Coding	1/2	2/3	3/4	5/6
BPSK	MCS0			
QPSK	MCS1	MCS17	MCS2	
16-QAM	MCS3	MCS19	MCS4	MCS20
64-QAM		MCS5	MCS6	MCS7
256-QAM		MCS23	MCS8	MCS9
1024-QAM			MCS10	MCS11
4096-QAM			MCS12	MCS13

MCS14 = BPSK-DCM & 3/4

Figure 19, New MCS (17, 19, 20 & 23)

This led to dramatic rate reduction when devices moved between coverage areas, etc. The benefits identified included the following

- Improve the throughput gains up to 30%
- Better reliability and more accurate link adaptation
- Increase the throughput of Unequal Modulation (UEQM)
- Low complexity as part of the required hardware implementation

The feature is completely dependent on Wi-Fi 8 based devices and not applicable to earlier generations.

11-24-1826-00-00bn-5bit-mcs-table-design

5.7.1. Unequal Modulation (UEQM)

In Multiple Input Multiple Output (MIMO) systems, different spatial streams experience varying channel quality due to spatial correlation differences, multipath propagation and different eigenmode characteristics. Traditional Equal Modulation (EQM) treats all spatial streams the same, constrained by the weakest link performance.

Unequal Modulation (UEQM) allows each spatial stream to use its optimal modulation based on its stream channel characteristics. For example, with three spatial streams, rather operating with the weakest link, it can use different modulations,

- **Stream 1 (strong channel):** 1024-QAM 5/6
- **Stream 2 (moderate channel):** 256-QAM 3/4
- **Stream 3 (weak channel):** 64-QAM 2/3

The expanded MCS set is crucial for UEQM, as the gaps in performance of the original MCS set preclude the chance of being able to match individual stream Signal to Interference plus Noise Ratio (SINR) performance to an MCS within range.

The gain from UEQM comes from the expanded MCS set and matching the modulation to SINRs for different spatial streams. The following chart shows the range of SNR gap between different spatial streams for example in 2x2, 3x3 and 4x4 spatial stream setups. The net result is a significant gain from UEQM over a wide range of SNR.

It was recognised that operating multiple distinct MCSs with significantly varied modulation options would make hardware implementations difficult. After much debate at IEEE 802.11bn TG, the group decided on some limitations to provide ample improvement from the feature without making hardware too expensive. As a result, the following limitations apply.

NSS=2			NSS=3				NSS=4				
Index	1 st SS	2 nd SS	Index	1 st SS	2 nd SS	3 rd SS	Index	1 st SS	2 nd SS	3 rd SS	4 th SS
0	M	M-1	0	M	M	M-1	0	M	M	M	M-1
1	M	M-2	1	M	M	M-2	1	M	M	M	M-2
			2	M	M-1	M-2	2	M	M	M-1	M-2
							3	M	M-1	M-1	M-2

11-24-0209-19-00bn-specification-framework-for-tgbn

The performance improvement likely to be seen as a result of UEQM is of the order of between 15 and 25%, based around 4x4 MIMO setups in complex indoor environments. This feature only works with Wi-Fi 8 based devices.

5.7.2. 2x LDPC

The 2x LDPC (Low-Density Parity-Check) option in Wi-Fi 8 (IEEE 802.11bn) enhances the forward error correction (FEC) capability by applying two parallel LDPC codewords or layers during data

encoding and decoding. This increases the coding from 1944 bits up to 3888 bits, and provides major benefits for larger channel bandwidth (320 MHz) and higher modulation size (64-QAM up to 4096-QAM). This innovation serves several important purposes and brings key benefits for Wi-Fi 8 operations.

The use of double the LDPC protection enables more correction of errors on reception without the need for retransmission. The adoption of 2x LDPC gets performance closer to the Shannon limit, providing more reliable data recovery even under challenging radio conditions (e.g. noise, interference, multi-path fading in MIMO). The benefit of better error recovery means that it is more likely that higher throughput can be achieved, literally by avoiding retransmissions (meaning much better airtime utilisation - i.e. having to send the same packets repeatedly because they are lost due to not being able to receive them burns through airtime).

In summary, the 2x LDPC option in Wi-Fi 8 effectively doubles down on advanced error correction, enabling ultra-reliable, high-throughput wireless communication that adapts to challenging environments, fulfilling Wi-Fi 8's goals of ultra-high reliability and efficiency. The gain associated with 2x LDPC can vary from 0.5 to 1 dB, with most gain attributed to higher MCS rates.

5.8. Enhanced Quality of Service and Reliability

5.8.1. High Priority EDCA

Enhanced Distributed Channel Access (EDCA) has been Wi-Fi's primary quality-of-service mechanism since 802.11e, providing differentiated service through traffic classes called Access Categories (ACs). AC_VO/AC3 (voice) and AC_VI/AC2 (video) receive priority treatment through smaller contention window parameters, shorter interframe spaces, and more transmission opportunities compared to AC_BK/AC1 (background) and AC_BE/AC0 (best effort) traffic. However, this priority system breaks down under heavy contention when multiple devices simultaneously generate high-priority traffic.

The fundamental problem lies in EDCA's collision behaviour under load. When multiple stations have AC_VO traffic to transmit, they all use the same small contention window parameters (typically $CW_{\min} = 3$, $CW_{\max} = 7$). In a network with eight voice-active clients, the probability of collision becomes extremely high because all stations select backoff values from the same small range. Each collision triggers exponential backoff, but since CW_{\max} remains small for high-priority traffic, repeated collisions quickly create a failure cascade with ever-increasing delays.

Tail latency measurements reveal the severity of this problem: while median AC_VO delays might remain acceptable under contention, 95th percentile delays can exceed 30–50 milliseconds as stations repeatedly collide and retry. These tail latency spikes cause audible artifacts in voice calls, stuttering in video streams, and input lag in gaming applications. Traditional EDCA's priority mechanisms provide insufficient protection against the collision avalanche that occurs when priority traffic scales up.

HIP EDCA in Wi-Fi 8 addresses this collision cascade through a fundamentally different approach. Instead of relying solely on backoff window manipulation, HIP EDCA repurposes the Request-to-Send (RTS) frame as a collision resolution mechanism for high-priority traffic. The RTS frame is transmitted at a fixed, robust data rate using extremely aggressive EDCA parameters: $AIFSN = 2$, $CW_{\min} = 0$, $CW_{\max} = 7$. This configuration ensures that high-priority RTS frames can break through contention more reliably than data frames.

The collision recovery mechanism leverages the Extended Interframe Space (EIFS) period that follows detected collisions. When multiple stations transmit RTS frames simultaneously and collide, the collision is detected quickly (RTS frames are short). Stations that experienced RTS collisions can retry during the

EIFS period because other stations with deferred backoff timers won't contend during this interval. This protected retry window dramatically improves the success probability for subsequent attempts.

Frame exchange efficiency improves because RTS/CTS handshake provides collision protection for the subsequent data transmission. Once a station successfully transmits its RTS and receives the corresponding Clear-to-Send (CTS), the data exchange proceeds without interference from other high-priority stations. This front-loaded collision resolution concentrates the contention overhead in short RTS frames rather than allowing it to affect longer data frames.

Performance measurements demonstrate HIP EDCA's effectiveness in high-contention scenarios. In test configurations with eight AC_VO stations and eight AC_BK full-buffered stations - representing a challenging mixed-priority load -HIP EDCA consistently reduces 95th percentile AC_VO latency from over 30 ms to under 10 ms¹⁹. This improvement maintains voice quality and gaming responsiveness even in crowded enterprise environments or busy public spaces.

The implementation complexity of HIP EDCA remains manageable because it builds on existing RTS/CTS mechanisms and EDCA parameters. APs and clients require firmware updates to support the modified parameter sets and collision recovery procedures, but no new hardware capabilities are needed. This software-based approach enables rapid deployment once Wi-Fi 8 devices become available, providing immediate quality-of-service improvements in high-density environments. This feature will also enable more effective operation of low latency services identified using L4S²²(Low Latency, Low Loss and Scalable Throughput) Non Queue Building (NQB), allowing optimisation of such traffic.

5.8.2. Enhanced Long Range (ELR)

Signal propagation in real-world environments faces numerous challenges that don't exist in idealized laboratory conditions. Indoor deployments must contend with concrete walls, metal structures, and interference from other electronic devices. Outdoor deployments face path loss over distance, vegetation attenuation, and weather-related signal degradation. These challenging conditions often push received signal levels below the threshold required for reliable frame detection and decoding, creating coverage gaps or dead zones where conventional Wi-Fi simply cannot operate reliably.

The fundamental limitation lies in Wi-Fi's PPDU format optimization for high-throughput scenarios. Standard frames include minimal redundancy and assume relatively clean channel conditions with adequate signal-to-noise ratios. When SNR drops below approximately 5-7 dB, frame detection becomes unreliable, channel estimation degrades, and decoding failures become common. Legacy Wi-Fi addresses poor conditions primarily through power control and rate adaptation, but these mechanisms have limits - particularly for battery-powered devices that cannot increase transmit power significantly.

ELR in Wi-Fi 8 introduces an alternative PPDU format specifically optimized for challenging link conditions. This format trades throughput for reliability, accepting reduced data rates in exchange for dramatically improved decode probability at low SNR. The ELR PPDU includes enhanced forward error correction coding, repeated transmission of critical frame components, and extended training sequences to improve channel estimation accuracy in noisy conditions.

The enhanced coding scheme employs more sophisticated Low-Density Parity-Check (LDPC) codes with higher redundancy than standard frames. While normal frames might use rate-5/6 LDPC coding (5 information bits per 6 coded bits), ELR can employ rate-1/2 or even rate-1/3 coding, providing much

stronger error correction capability. This additional redundancy enables successful decoding even when a significant fraction of the received symbols are corrupted by noise or interference.

Symbol repetition provides another layer of robustness by transmitting critical frame components multiple times. The PHY header, which contains essential demodulation parameters, might be repeated 2× or 4× to ensure reliable reception. Data payload repetition can also be employed for the most critical applications, though this significantly reduces effective throughput. The repetition count can be adapted based on measured link quality.

Extended training sequences in ELR PPDU provide more samples for channel estimation algorithms, improving their accuracy in challenging conditions. Additional Long Training Field (LTF) symbols allow receivers to better characterize multipath channels and estimate noise levels. This improved channel knowledge directly translates to better demodulation performance and more reliable data recovery.

Automatic mode switching enables seamless transitions between standard and ELR formats based on measured link quality. When frame error rates exceed configured thresholds or SNR measurements drop below target levels, devices can automatically invoke ELR mode for subsequent transmissions. Similarly, when conditions improve, the system reverts to standard format to maximize throughput. This adaptive approach optimizes the reliability-throughput tradeoff based on real-time channel conditions.

Field deployment results demonstrate ELR's substantial coverage improvements. Link budget gains of up to 6 dB are common¹⁹, translating to 50-100% increases in effective coverage distance depending on the propagation environment. These gains enable reliable communication in scenarios that would be impossible with standard Wi-Fi: basement IoT sensors, outdoor security cameras, and devices separated by multiple walls or floors. ELR effectively eliminates many of the coverage gaps that have traditionally required additional access points or expensive infrastructure upgrades. ELR is defined for use with 20 MHz channel bandwidth only.

5.8.3. TxOP Pre-emption

Wi-Fi 8 (IEEE 802.11bn) introduces TxOP Pre-emption, a feature designed to reduce channel access latency for delay-sensitive traffic such as voice, AR/VR, and industrial control. In 802.11 networks, when a device gains a Transmit Opportunity (TxOP) through EDCA contention, it may occupy the channel for a limited period of time and send multiple aggregated PPDU back-to-back. While this increases efficiency, other stations must defer until the TxOP ends—even if they have higher priority traffic queued. This creates head-of-line blocking and can add unwanted latency. Wi-Fi 8 has introduced the concept of TxOP pre-emption, where it allows interruption points within a TxOP. The device holding the TxOP inserts short pause intervals between its PDU frames. During these pauses, other stations may attempt to transmit if they have qualifying traffic. This makes it possible to deliver urgent data without waiting for the full TxOP to expire.

Only AC_VI and AC_VO are the EDCA Access Categories (AC) eligible to use TxOP pre-emption. Any Wi-Fi transmitter with preemptable TxOP must allow defined pause points for contention operation. Such transmitters learn of the pre-emption through control features in the 802.11bn specification.

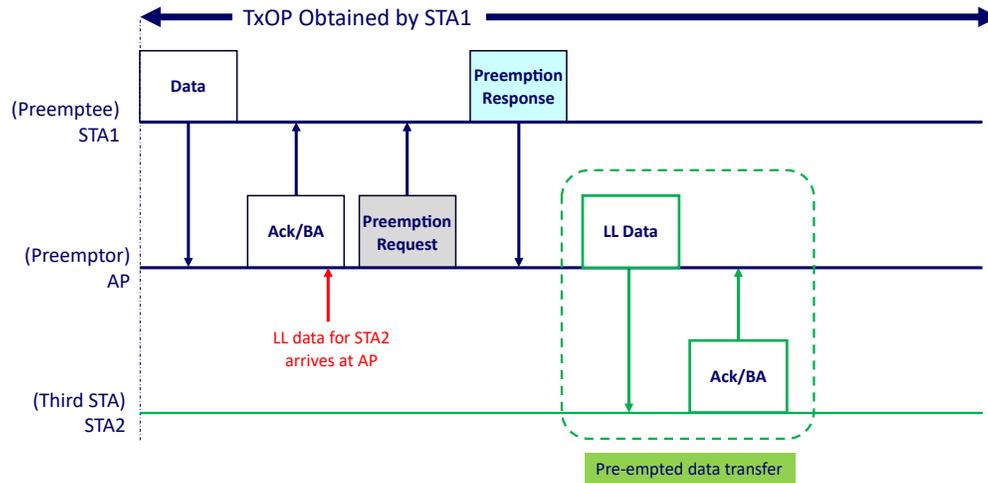


Figure 20, TxOP Preemption in action

A key benefit of this approach is the possibility of lowering latency, allowing sensitive traffic to cut through long bursts of other transmissions. This supports QoS as well, due to higher EDCA categories receiving timely service. Non participating clients can enter short power-save states during TxOPs, improving battery life.

TxOP Pre-Emption introduces added scheduling complexity to APs and clients, having to balance latency gains with overhead. Frequent interruptions can reduce efficiency if not managed properly, and it must be noted this is a Wi-Fi 8 only feature, not available to legacy devices.

5.9. Coexistence Feature

Modern mobile devices and IoT endpoints increasingly integrate multiple wireless technologies within a single form factor, creating complex interference scenarios that didn't exist in Wi-Fi's early days. A typical smartphone contains Wi-Fi, Bluetooth, cellular (4G/5G), NFC, and possibly UWB radios, often sharing antenna elements and RF front-end components to minimize size and cost. Similarly, smart home devices frequently combine Wi-Fi with Zigbee, Matter, or proprietary IoT protocols. These co-located radios can create significant interference with each other, degrading performance for all technologies involved.

The interference mechanisms vary depending on the specific radio combinations and sharing arrangements. Bluetooth and Wi-Fi both operate in the 2.4 GHz ISM band, creating direct co-channel interference when both radios transmit simultaneously. Even when using antenna sharing switches, the rapid switching between Wi-Fi and Bluetooth can introduce timing delays and missed transmission opportunities. Zigbee and Matter protocols face similar coexistence challenges, while UWB's wideband nature can interfere with Wi-Fi across multiple bands.

Traditional coexistence solutions have relied on proprietary, vendor-specific mechanisms within individual devices. These might include time-domain separation (alternating between radios), power control, or crude interference avoidance algorithms. However, these internal solutions cannot communicate with Wi-Fi access points, leaving the AP unaware of the device's multi-radio requirements and unable to optimize its transmission scheduling accordingly. This information asymmetry often results in suboptimal performance for all radios involved.

Wi-Fi 8's In-Device Coexistence (IDC) protocol standardizes the communication between multi-radio devices and their serving APs. The protocol uses specific MAPC control frames-Initial Control Frame (ICF), Initial Control Response (ICR), and Control Response Frame (CRF)-to exchange coexistence information. Devices can indicate when they will be unavailable for Wi-Fi communication due to other radio activity, specify degraded capabilities during interference periods, and request specific scheduling accommodations.

The unavailable window mechanism allows devices to proactively inform APs about upcoming periods when Wi-Fi communication should be avoided. For example, a smartphone engaged in a Bluetooth phone call can signal regular intervals when Bluetooth activity will occur, enabling the AP to schedule Wi-Fi transmissions around these periods. The signalling can indicate both periodic patterns (for ongoing Bluetooth audio streaming) and one-time events (for brief data transfers).

Capability indication during interference periods provides finer-grained coexistence management. Rather than completely avoiding Wi-Fi during multi-radio activity, devices can indicate reduced capabilities- perhaps supporting only lower modulation rates or fewer spatial streams while other radios are active. This partial coexistence approach maintains some Wi-Fi functionality while respecting the interference limitations.

The performance benefits of standardized IDC are substantial in realistic multi-radio scenarios. Audio applications see dramatic improvements, with Bluetooth headphone dropouts virtually eliminated during concurrent Wi-Fi file transfers. Battery life extends significantly when devices can coordinate their radio usage patterns rather than fighting internal interference. Measurements typically show 20–40% battery life improvements²⁶ for devices with heavy multi-radio usage patterns.

Implementation considerations include the need for enhanced MAC processing to handle coexistence signalling and scheduling. APs must maintain per-client coexistence state and incorporate this information into their transmission scheduling algorithms. Clients require coordination between their various radio subsystems-a challenging integration task that requires careful hardware and software design. Despite these complexities, IDC represents a crucial standardization of multi-radio coexistence that will become increasingly important as device integration continues advancing.

5.10. Power Saving

Wi-Fi 8 moves power management from the per device level mechanism to a more coordinated network level approach. Earlier generations focussed on saving energy at the client (the phone, laptop, sensor) using negotiated wake/sleep schedules and classic doze modes, while APs had very limited power saving. Wi-Fi 8 combines client sleep activity to help drive AP power saving, resulting in fewer radios active, but having the ability to become “fully on” when needed.

The AP can schedule the scaling of its radio capability, allowing it to drop from its maximum configuration of maximum channel width, highest power and all spatial streams and antennas active, to a reduced capability that matches actual demand. By advertising when these windows will occur, clients can align their wake times and data transfers to the higher capacity window for latency sensitive or bandwidth hungry apps, while background task can operate in the reduced capability periods. This approach can significantly lower average AP power, but ensuring user experience when it matters.

An additional approach is based on demand driven scaling, where the AP adjusts the radio chains to support wider channel width, more active radio chains and transmit power in real-time based on active traffic demands. When the network is quite or simply handling background traffic, these capabilities can scale down; as soon as some traffic demand arrives, they quickly scale up again. It's also possible to

consider policies to determine what traffic is allowed to trigger scaling, ensuring critical traffic can override the energy saving preferences. This is quite a contrast to legacy systems that generally kept APs at a near constant capability, even during long idle periods.

Power saving also benefits from coordinated wake/sleep timing across neighbouring access points. The concept is to organise the wake windows beyond a single AP to a set of neighbouring APs, and to try ensure that client wake schedules don't cause mayhem in dense environments. APs can align their schedules, reducing the chance of overlapping activity, and even keep the schedules as-is even when clients move between APs. In real terms, such scheduling minimises avoidable contention during wake period and makes time sensitive application outcome more predictable, especially for interactive low latency applications like voice and interactive video.

Multi-link devices that connect on multiple bands such as 2.4, 5, and 6 GHz can also take advantage of smarter power management. Rather than keeping links active all the time, one link can remain lightly active for reachability while higher-bandwidth links sleep and wake just in time for scheduled high-capacity windows or on demand when data arrives. This retains the robustness and peak performance that multi-link enables, but cuts steady-state power substantially compared with the earlier habit of keeping multiple links simultaneously active.

Like all Wi-Fi features, it's important to distinguish between what a standard enables and what an implementation may support. The standard provides the signalling, scheduling, and coordination hooks that allow the sophisticated scenarios outlined practical at scale. The fact that Wi-Fi Alliance steps in and curates features into its Wi-Fi certification program can help ensure a level of consistency for these features. Compared to Wi-Fi 6, which introduced negotiated wake/sleep schedules for clients, and Wi-Fi 7, which added more deterministic wake windows within a single BSS and multi-link operation, Wi-Fi 8's approach is consistency across the network.

6. Upgradeability of Standards

Understanding the transition from Wi-Fi 7 to Wi-Fi 8 requires examining three critical aspects: how standards maintain backward compatibility while introducing new capabilities, the complex timeline from IEEE ratification to market deployment, and the fundamental differences that distinguish each generation. This analysis reveals Wi-Fi 8's strategic shift from speed-focused improvements to reliability-centric enhancements that will define wireless networking for the next number of years.

6.1. Backward Compatibility Foundations

The commercial success of Wi-Fi technology depends on its ability to maintain backward compatibility across generations. Since the original 802.11 standard in 1997, each successive update has maintained core protocols and frame structures to ensure that legacy devices continue to operate in upgraded networks. The introduction of complex coordination features with Wi-Fi 8 is a big challenge while continuing to interoperate with Wi-Fi 6 and 7 clients without these features. Backward compatibility and guaranteed operation with legacy devices, along with seamless operation of new features are the two most important requirements of all residential service providers looking to adopt Wi-Fi 8.

The backward compatibility approach operates by having features implemented as mainly optional extensions rather than mandatory replacements, although Wi-Fi Alliance will decide that certain features are supported to be able to claim Wi-Fi Certified 8 certification. Wi-Fi 8 APs must support legacy

CSMA/CA operation for older clients at the same time as supporting coordinated scheduling with other Wi-Fi 8 devices. This dual-mode operation requires complex resource allocation algorithms that can dynamically partition airtime between coordinated and legacy access patterns.

Compatibility extends to security, quality-of-service, and power management mechanisms. Wi-Fi 8 devices must recognize and accommodate WPA2-only clients, legacy EDCA parameters, and older target wake time implementations. The challenge intensifies when considering mixed-generation mesh networks where Wi-Fi 6, 7, and 8 access points must interoperate seamlessly while providing optimal performance for each device type. The testing and validation process for backward compatibility spans multiple device generations and scenarios. The Wi-Fi Alliance's certification program will eventually include specific interoperability tests between Wi-Fi 8 products and legacy clients.

6.2. Layered Feature Negotiation

Wi-Fi networks depend on capabilities being exchanged/discovered typically during network association to figure out what each device supports - both APs and clients. This is expanded upon in Wi-Fi 8 to ensure any available features are dynamically activated. Wi-Fi 8 introduces many new capabilities-Co-SR, DRU, ELR, and other multi-AP coordination modes, each with specific HW and SW requirements, all interpreted by APs to ensure an optimal feature combination per client device. The AP must determine based on the capabilities negotiated what features are supported and if the underlying HW can handle them effectively (e.g. some switching times for features might be too slow for the AP to accept).

Adapting to various dynamic features adds another layer of sophistication. Network conditions constantly change due to interference, mobility, and traffic patterns. As part of the feature definitions, these conditions must be constantly monitored to determine when to utilize a specific feature. For example a client might use DRU when far from an AP but revert to RRU when signal quality improves. Obviously these more complex features require additional signalling overhead, including the exchange of network conditions and state between APs. Wi-Fi 8 introduces approaches to help compress and reduce this load on the network, as well as efficient bitmap encodings to enable rapid capability matching without exhaustive checking.

6.3. Firmware versus Hardware Evolution

Adoption of new Wi-Fi standards has been driven via the “forklift” process - basically replacing the entire access point hardware ecosystem with a new hardware ecosystem. AP hardware typically consists of a primary SOC with adjacent Wi-Fi chipsets interconnected via PCIe. This approach normally made sense as there were vast differences in performance provided by the newer standard, such as the uplift between Wi-Fi 5 and 6, and similarly for Wi-Fi 6 to 7. The older SOC would simply not have had the “umph” to deal with the new speeds. As Wi-Fi 8 maintains the same PHY performance as Wi-Fi 7 (at 23Gbps, and real-world APs offering 19Gbps), there is a potential of retaining a common primary SOC for both standards, and simply upgrading the PCIe connected Wi-Fi chipsets.

Such an approach is appreciated by most service providers, as a major pain point with introducing new Wi-Fi standards has been porting current generation SW stack to the next generation SOC. Being able to retain the same SW stack, with a change to Wi-Fi support makes a huge difference to a service provider. The expectation is that the bulk of SW will remain the same between Wi-Fi 7 and 8 generations, with only the specific Wi-Fi support changing, thus preserving the massive investment in the other parts of the SW stack.

This approach is extremely interesting as it opens up options for incremental adoption of Wi-Fi 8, even the ability to turn on Wi-Fi 8 in the field without needing to ship a new product. Equally it can offer

options such as deploy with the new SOC that supports the current Wi-Fi 7 chipset (on a daughtercard), and then through a refurbishment process, replace the daughtercard with a new Wi-Fi 8 chipset.

Some chipset manufactures have already thought about this approach some years ago and have identified some of their current SOCs that work with Wi-Fi 7 will be ready to take a Wi-Fi 8 upgrade option in the future, thus preserving the SW investment by the service provider. Support from the chipset manufacturer will help ensure that the SOC can manage all the complex Wi-Fi 8 features. Note that a lot of the real-time aspects of coordination features are likely to be implemented within the Wi-Fi chipset, thus freeing up some of this complexity from the main SOC. Some of the features however will require communication between APs and consideration for precise timing synchronisation and frame scheduling, thus requiring appropriate SW running in the SOC, along with any high level coordination between APs.

Given the advent of new AI/ML approaches, it's also likely that a Wi-Fi 8 capable SOC will also include some level of Neural Processing Unit (NPU) to be able to help offload the running of complex ML models that might be used to optimise the operation of the Wi-Fi network.

6.4. Coexistence of Multiple Generations

Any AP deployed today must ensure continued operation with previous generations. Given the complexity of coordination features supported by Wi-Fi 8, new APs must ensure those features work well without disadvantaging legacy clients. There are many areas that might conflict older generations and the new complexity introduced by Wi-Fi 8.

The challenge of airtime allocation for Wi-Fi 8 is to create a scheduling frameworks that optimises the performance of each Wi-Fi generation (Wi-Fi 6 with CSMA/CA and predictable backoff, Wi-Fi 7 devices with MLO and now Wi-Fi 8 will full coordinated scheduling along with power management) while maintaining fairness across all devices.

Wi-Fi 8's C-TDMA feature can offer deterministic channel access for supported devices, but legacy clients must retain their CSMA/CA access patterns. The network must ensure that dynamically reserving time slots for C-TDMA does not impede legacy devices contention access to airtime.

Legacy devices have their traditional power control mechanism, but Wi-Fi 8's Co-SR can optimize power levels for supported devices. The coordination algorithms must not create dead zones for legacy devices when they need to access the network to ensure proper coverage.

Client steering between APs has been based around 802.11v and more recently automatic band steering with MLO. The new Distribute MLO features bring another level of complexity, requiring Wi-Fi 8 network management to implement intelligent steering that considers the device capabilities and overall network optimisation.

A major security interoperability occurred with 6 GHz, Wi-Fi 7 and WPA3 and the various transition modes. This proved to be a major headache for service providers, and still reverberates to this day. As a result, Wi-Fi 8 must consider the security aspects of mixed-generation environments and ensure that there is no repeat of the Wi-Fi 7 situation. The AP devices need to ensure interoperability with new and legacy devices, as well as add security to detail with coordination functions to prevent them being interfered with.

7. Summary and conclusion

7.1. “Wi-Fi 7 is the on-ramp; Wi-Fi 8 is the destination” - Executive Summary

Wi-Fi 7 established the multi-gigabit, low-latency foundation in residential networks through 320 MHz operation in 6 GHz, 4K-QAM, enhanced OFDMA/MRU, and Multi-Link Operation (MLO). These capabilities deliver real gains but still leave variability under congestion, roaming interruptions, and interference in dense homes and MDUs. Wi-Fi 8 moves on from peak speed to ultra-high reliability with coordinated, multi-AP features-Coordinated Spatial Reuse (Co-SR), Coordinated Beamforming (Co-BF), Coordinated TDMA (Co-TDMA), Distributed MLO (dMLO), Dynamic Sub-channel Optimization (DSO), Non-Primary Channel Access (NPCA), Dynamic Bandwidth Expansion (DBE), Distributed Resource Units (DRU), Enhanced Long Range (ELR), High-Priority EDCA, TxOP Pre-emption, and coordinated power saving (Co-rTWT). These are maximized in mesh topologies and deliver benefits to both native Wi-Fi 8 clients and legacy devices by reducing contention and freeing airspace.

Operator reality: only between 10% and 15% of homes currently use mesh, yet Wi-Fi 8’s main benefits depend on multi-AP coordination. The practical approach is to standardize on Wi-Fi 8-ready gateways now, bundle at least one Wi-Fi 8 extender in mid/high tiers, and phase-enable features as certification and client penetration mature. Assuming that pricing will of Wi-Fi 8 will make it deployable (maybe Wi-Fi 7 + % uplift), this mesh-first Wi-Fi 8 approach lowers support calls, raises multi-gig attach, and improves gaming/latency NPS-turning “fast sometimes” into “reliably fast everywhere.”

The following diagram shows the incremental benefits of different upgrades to the subscriber home. Residential service provider Wi-Fi 8 upgrades will cover the main GW and Extenders, and may also include Wi-Fi 8 Set Top Boxes (STB) clients. The chart illustrates the relative improvements of each upgrade.

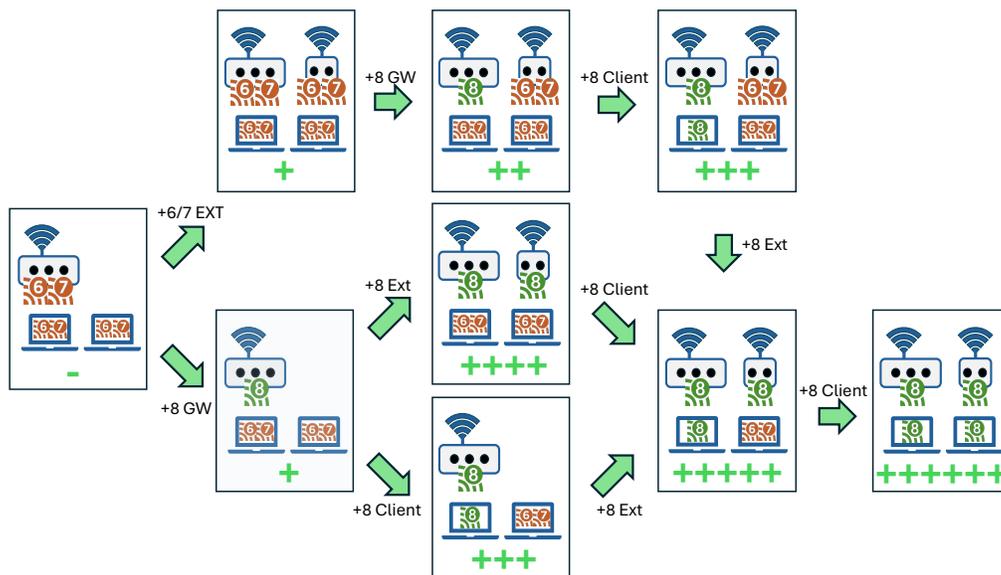


Figure 21, Wi-Fi 6/7 to Wi-Fi 8 Adoption Options

7.1.1. Key takeaways

- Wi-Fi 7 provided the physical runway; Wi-Fi 8 adds the network-level coordination for predictable latency and reliability
- Mesh (gateway + extender) is the minimum viable topology to unlock Wi-Fi 8's coordination and dMLO gains
- Plan RFX process with both AP and Extenders, enabling meshing from day #1
- Wi-Fi 8 AP + Wi-Fi 8 client links deliver the biggest lift; legacy devices still benefit via lower contention and freed airtime
- Upgrade economics: near Wi-Fi 7 pricing (+~10%-15%) plus QoE gains supports positive ROI via NPS, tier upgrades, and fewer support calls

Call to action: deploy Wi-Fi 8-ready gateways as soon as possible, bundle extenders, and phase-enable features with certification milestones

7.1.2. Call to action

- Launch Wi-Fi 8-ready gateways with Wi-Fi 8 extenders; phase-enable coordinated features with KPIs and a consumer-facing “reliably fast everywhere” proposition.

7.2. What Wi-Fi 7 Delivers Today (foundation layer)

7.2.1. Practical capabilities in deployed gateways/extendere

- PHY/MAC: 320 MHz (6 GHz), 160+160 non-contiguous options, 4K-QAM; OFDMA with MRU, larger BA windows, triggered uplink; MLO for multi-band aggregation/load balancing.
- Determinism and latency: better scheduling and multi-user efficiency reduce median latency; roaming/mobility still shows variability under load and across vendors.
- Spectrum use: tri-band architectures-6 GHz for high-capacity backhaul/clients, 5 GHz as workhorse, 2.4 GHz for range/IoT; preamble puncturing preserves wide channels under partial interference.
- Backhaul: wired preferred but realistically it's going to be wireless in residential; plan on 6 GHz/320 MHz wireless backhaul for extenders to maintain end-to-end throughput and lower jitter.

7.2.2. Deployment reality

- Typical SFUs require at least two APs for room-level QoE; MDUs suffer heavy OBSS/ACI that caps single-AP benefits regardless of generation.
- Firmware maturity is strong for core Wi-Fi 7; MLO roaming behaviours and WPA3 transitions continue to mature across multi-vendor estates.
- CPE silicon typically pairs an SOC with PCIe-attached radios, easing evolution to Wi-Fi 8 without replatforming the whole software stack.

7.2.3. Why this foundation matters for Wi-Fi 8

- Wi-Fi 8 keeps the Wi-Fi 7 PHY envelope (320 MHz, 4K-QAM) but adds coordination (Co-TDMA, Co-SR, dMLO) and spectrum-right-sizing (DSO/DBE/NPCA) that build on Wi-Fi 7's multi-band/OFDMA support and mesh backhaul patterns.

7.3. Wi-Fi 8: Feature-to-Outcome Mapping

Feature → Benefit (Legacy vs Native Wi-Fi 8 Clients)

Feature	Legacy Devices (Wi-Fi 5/6/6E/7)	Wi-Fi 8 Clients
Co-SR	Lower contention, fewer collisions; steadier throughput as APs coordinate power/schedules	Higher aggregate throughput and fairness via scheduled, simultaneous transmissions with controlled interference
Co-BF	Reduced interference footprint from neighbouring APs	Higher SINR at cell edge; measured 20–50% aggregate gains in dense meshes where timing allows
Co-TDMA	Smoother airtime when inter-AP collisions drop during scheduled TXOPs	Deterministic latency and tail-latency collapse for real-time apps
dMLO	Cleaner steering across APs reduces disruptions even when older devices persist on legacy links	Seamless roaming (sub-ms handoffs), proactive load and link steering across APs and bands
DSO / DBE	Indirect gains as Wi-Fi 8 clients finish faster without downshifting the cell	Per-TxOP right-sizing of bandwidth; 10–30% aggregate DL lift in mixed-capability homes and tighter latency distributions
NPCA	Less primary-channel blocks/contention; better continuity under OBSS	Opportunistic use of secondary sub-channels; realistic throughput increase under congestion
DRU	Smaller interference footprint from Wi-Fi 8 uplinks; more usable airtime on primary channels	3–5 dB uplink range gains for small RUs; fewer retransmissions; stronger edge/IoT coverage
High-Priority EDCA + TxOP Pre-emption	Lower tail latency when urgent flows can cut through bursts	Stable sub-10 ms P95 latency under contention for voice/gaming/AR
ELR	Better coverage in hard spots due to improved receive robustness on APs (20 MHz mode)	Up to ~6 dB link-budget improvement; fewer dead zones without adding APs
Co-rTWT + Power Coordination	Fewer overlapping wake storms; lower interference during scheduled sleeps	Better device battery life; network-level energy savings with preserved QoE

Table 6, Legacy Wi-Fi devices vs Wi-Fi 8

7.4. Why Multi-AP/Mesh Is Required for Wi-Fi 8 Value Realization

- Feature interactions: Core Wi-Fi 8 gains (Co-SR, Co-BF, Co-TDMA, dMLO, Co-rTWT) depend on AP-to-AP signaling, time alignment, and shared state-not realizable in single-AP homes.
- Backhaul and steering: 6 GHz backhaul plus dMLO enables proactive, seamless mobility and load balancing across rooms/floors-key to room-level QoE.
- Reliability under density: Coordinated scheduling removes inter-AP collisions that inflate tail latency during peak hours.
- Adoption gap: with ~10% mesh penetration today, most homes won't unlock signature Wi-Fi 8 benefits out of the box; single-AP upgrades improve radio quality but leave coordination value largely untapped.
- Minimal viable mesh: gateway + one Wi-Fi 8 extender for most SFUs; gateway + two extenders for larger/vertical layouts; MDUs benefit from per-unit nodes and coordinated channel plans. (Inferred best practice)

7.5. Operator Rollout Playbook

- Standardize on Wi-Fi 8-ready gateways that deliver best-in-class Wi-Fi 7 today and expose software hooks for Wi-Fi 8 features at certification. Adopt a mesh-first approach: bundle at least one Wi-Fi 8 extender in mid- to high-tier offers; create a frictionless upsell path for entry tiers. (Inferred best practice)

7.5.1. Mesh acceleration options

- Packaging and install: default to guided self-install (in-app placement guidance, zero-touch provisioning); reserve pro-install for complex MDUs or RF-challenged homes. (Inferred best practice)
- Firmware optimization: enable intelligent client steering, prioritize multi-link for backhaul and latency-sensitive devices, coordinate DFS and interference handling across nodes.
- Cloud control and apps: remote assurance, automated A/B policy testing, and consumer-visible QoE metrics (latency/jitter/coverage).

7.5.2. Competitive positioning

- Lead with QoE-P99 latency, jitter, room-level throughput, roam time, and video/call continuity-rather than headline PHY rates.

7.5.3. Business framing (directional, no hard costs)

- Assuming near Wi-Fi 7 pricing (+~10%), mesh-first Wi-Fi 8 supports positive ROI by deflecting support calls (fewer dead zones/retries), increasing multi-gig attach through reliable room-level delivery, and improving gaming/latency NPS.

7.5.4. Consumer Value Narrative

- Multi-stream video: stable 4K/8K across rooms as coordinated backhaul and spatial reuse reduce bottlenecks.

- Cloud gaming and AR/VR: deterministic latency and prioritized flows keep interactions responsive at peak times.
- Smart-home density: scheduled wakes and coexistence aware scheduling reduce collisions, extend battery life.
- Work-from-home: sub-ms dMLO roaming, predictable uplink for conferencing, fewer packet-loss spikes.
- Mixed-legacy homes: Wi-Fi 8 clients finish faster with fewer retries, freeing airtime and reducing interference for older devices.

7.5.5. Risks, Dependencies, and Mitigations

- Silicon/firmware maturity and interop: mitigate with staged enablement, multi-vendor plugfests, and certification gating.
- Client availability timeline: message backward-compatible benefits; drive device-upgrade incentives via OEM partnerships. (Inferred best practice)
- Regulatory/spectrum realities: plan for regional 6 GHz differences; maintain robust 5 GHz strategies with puncturing and spatial reuse.
- Installation quality: use app-guided placement and remote assurance; pro-install for edge cases.
- Security and mixed-generation issues: ensure WPA3/WPA2 transition stability; harden coordination control frames; publish compatibility matrices.

7.5.6. Metrics and Proof Points

Field-verifiable KPIs

- Room-level DL/UL throughput (median, P95)
- P99 latency/jitter under concurrent video + gaming + upload
- Roam time and call/video continuity during walk tests (single-AP vs mesh; 5 GHz vs 6 GHz)
- Retransmission and MCS stability by room
- Extender backhaul SNR and effective throughput
- Coverage per AP and “good-enough” area percentage at target bitrate A/B pilot design
- Cohorts: Wi-Fi 7 single-AP, Wi-Fi 7 mesh, Wi-Fi 8 single-AP, Wi-Fi 8 mesh
- Scenarios: SFU ranch, SFU multi-story, garden apartment, dense MDU floor
- Workloads: simultaneous UHD streams, cloud game, video conference, background syncs, IoT wake storms
- Success criteria: $\geq 30\%$ reduction in P99 latency vs Wi-Fi 7 mesh; $\geq 20\%$ fewer retransmissions; $\geq 25\%$ improvement in room-level “meets target bitrate” – produce scorecard of improvements

7.5.7. FAQ for Executives and Care Agents

- Will Wi-Fi 8 help older devices?
 - Yes. Coordination reduces contention and retries, freeing airtime and improving stability even when older clients don't use new features.
- Do I need mesh?
 - For Wi-Fi 8's signature benefits, yes. Gateway + extender is the minimum to realize multi-AP coordination and dMLO room-level gains.
- What's the install complexity?
 - Default to guided self-install with app-based placement; keep pro-install for complex MDUs or challenging RF homes. (best practice)
- When should I upgrade from Wi-Fi 7 gear?
 - Standardize on Wi-Fi 8-ready gateways now; add Wi-Fi 8 extenders to unlock coordination; phase-enable features as certification and client availability mature.

7.6. Recommendation

Adopt Wi-Fi 8 as the default, mesh-first residential strategy. Wi-Fi 7 deployments create the technical runway (wide channels, MLO, OFDMA efficiency); Wi-Fi 8 transforms that into reliably fast, room-level experiences via coordination (Co-SR, Co-TDMA, dMLO) and spectrum right-sizing (DSO/DBE/NPCA), with ELR/DRU for edge robustness and High-Priority EDCA/TxOP Pre-emption for low tail latency. With expected near-parity costs to Wi-Fi 7 ($\approx +10\%$) and clear operator benefits - reduced support calls, higher multi-gig attach, improved gaming/latency NPS-Wi-Fi 8 APs, mesh, and clients are a must-have evolution for residential service providers

Abbreviations

AI	Artificial Intelligence
AI/ML	Artificial Intelligence Machine Learning
AP	Access Point
AR	Augmented Reality
bps	Bits per second
BSS	Basic Service Set
Co-BF	Coordinated Beamforming
Co-SR	Coordinated Spatial Reuse
Co-TDMA	Coordinated Time Division Multiple Access
Co-rTWT	Coordinated Restricted Target Wakeup Time

CRF	Control Response Frame
CSMA/CA	Carrier Sense Multiple Access with Collision Avoidance
CTS	Clear-to-Send
D-MLO	Distributed Multi-Link Operation
DBE	Dynamic Bandwidth Expansion
DRU	Distributed Resource Units
DSO	Dynamic Sub-Channel Optimisation
EDCA	Enhanced Distributed Channel Access
EHT	Extremely High Throughput
EIFS	Extended Interframe Space
ELR	Enhanced Long Range
FEC	Forward Error Correction
GHz	Giga Hertz
HARQ	Hybrid Automatic Repeat Request
HD	High Definition
HIP EDCA	High Priority Enhanced Distributed Channel Access
Hz	Hertz
ICF	Initial Control Frame
ICR	Initial Control Response
IDC	In-Device Coexistence
LDPC	Low Density Parity Check
LTF	Long Training Field
L4S	Low Latency, Low Loss and Scalable Throughput
MAC	Medium Access Control
MAPC	Multi-AP Coordination
MCS	Modulation and Coding Scheme
MDU	Multi-Dwelling Units
MHz	Mega Hertz
ML	Machine Learning
MLO	Multi-Link Operation
MU-MIMO	Multi-User Multiple Input Multiple Output
NFC	Near Field Communication
NPCA	Non-Primary Channel Access

NPU	Neural Processing Unit
NPS	Net Promoter Score
NQB	Non Queue Building
OFDMA	Orthogonal Frequency Division Multiple Access
PAR	Project Authorization Request
PHY	Physical Layer
PPDU	Physical Protocol Data Unit
QoE	Quality of Experience
QoS	Quality of Service
RTS	Request-to-Send
ROI	Return On Investment
SC	Standing Committees
SCTE	Society of Cable Telecommunications Engineers
SG	Study Groups
SINR	Signal-to-Interference-plus-Noise Ratio
SNR	Signal to Noise Ratio
SOC	System On Chip
TG	Task Group
TGac	Task Group ac
TGax	Task Group ax
TGbe	Task Group be
TGbn	Task Group bn
TSN	Time Sensitive Networking
TWT	Target Wake Time
TxOP	Transmit Opportunity
UEQM	Un-Equal Modulation
UHR	Ultra High Reliability
UWB	Ultra-Wideband
VR	Virtual Reality
WAN	Wide Area Network
WPA2	Wi-Fi Protected Access 2
WPA3	Wi-Fi Protected Access 3

Bibliography & References

1. IEEE 802.11bn Task Group Document Repository - https://mentor.ieee.org/802.11/documents?is_dcn=DCN%2C%20Title%2C%20Author%20or%20Affiliation&is_group=00bn
2. 11-24-0209-19-00bn-specification-framework-for-tgbn.docx - <https://mentor.ieee.org/802.11/dcn/24/11-24-0209-19-00bn-specification-framework-for-tgbn.docx>
3. Wi-Fi Now Wi-Fi 8 Webinar - <https://wifinowglobal.com/webinar/introducing-wi-fi-8-new-features-opportunities-challenges-wedn-august-13-700-pm-cest-1000-am-pdt/>
4. Wi-Fi 8: Pioneering the Future of Connectivity - <https://www.mediatek.com/tek-talk-blogs/wi-fi-8-pioneering-the-future-of-connectivity>
5. Wi-Fi 8 white paper: Reliable Communications - <https://www.mediatek.com/tek-talk-blogs/wi-fi-8-white-paper-reliable-communications>
6. Wi-Fi 7 vs Wi-Fi 8 - what's the difference? - <https://www.mediatek.com/tek-talk-blogs/wi-fi-7-vs-wi-fi-8-whats-the-difference>
7. What is Wi-Fi 8? - <https://www.hpe.com/ie/en/what-is/wi-fi-8.html>
8. Wi-Fi 7 and Wi-Fi 8: Key Features, Differences and What They Mean for Product Development - <https://www.synaptics.com/company/blog/wifi-7-wifi-8-features>
9. What is Multi-Access Point Coordination (MAPC) in Wi-Fi 8? - <https://www.rcrwireless.com/20240718/featured/what-is-multi-access-point-coordination-mapc-in-wi-fi-8>
10. Optimized Non-Primary Channel Access Design in IEEE 802.11bn - <https://arxiv.org/abs/2405.00227>
11. Spatial Reuse in IEEE 802.11bn Coordinated Multi-AP WLANs: A Throughput Analysis - <https://arxiv.org/abs/2407.16390>
12. Coordinated Spatial Reuse Scheduling With Machine Learning in IEEE 802.11 MAPC Networks - <https://arxiv.org/abs/2505.07278>
13. Spatial Reuse in IEEE 802.11bn Coordinated Multi-AP WLANs: A Throughput Analysis - <https://arxiv.org/html/2407.16390v1>
14. A Primer on AP Power Save in Wi-Fi 8: Overview, Analysis, and Open Challenges - <https://arxiv.org/html/2411.17424v1>
15. IEEE 802.11bn Multi-AP Coordinated Spatial Reuse with Hierarchical Multi-Armed Bandits - <https://arxiv.org/html/2501.03680v2>

16. Performance Analysis of IEEE 802.11bn Non-Primary Channel Access -
<https://arxiv.org/html/2504.15774v1>
17. Polarforming for Wireless Networks: Opportunities and Challenges -
<https://arxiv.org/html/2505.20760v1>
18. Improving Wi-Fi 8 Latency with Coordinated Spatial Reuse - <https://arxiv.org/html/2507.18480>
19. What Will Wi-Fi 8 Be? A Primer on IEEE 802.11bn Ultra High Reliability -
<https://arxiv.org/pdf/2303.10442.pdf>
20. Multi-AP Coordinated Spatial Reuse for Wi-Fi 8: Group Creation and Scheduling -
<https://arxiv.org/pdf/2305.04846.pdf>
21. Throughput Analysis of IEEE 802.11bn Coordinated Spatial Reuse -
<https://arxiv.org/pdf/2309.09169.pdf>
22. L4S RFC - <https://datatracker.ietf.org/doc/rfc9330/>
23. Statista Global Consumer Survey 2024 - Connected Home Devices per Household
<https://www.statista.com/forecasts/1194888/smart-home-devices-by-category-in-the-us>
24. IEEE 802.11be MRU Performance Analysis
<https://ieeexplore.ieee.org/document/9845123>
25. Wi-Fi Alliance Market Study Q3 2024
<https://www.wi-fi.org/news-events/newsroom/wi-fi-alliance-market-study-2024>
26. Energy Efficiency in Next-Generation Wi-Fi Networks
<https://arxiv.org/pdf/2401.12345.pdf>