



OPEN
Compute Project



WHITE PAPER: SUSTAINING MULTICAST CONTINUITY FOR ROAMING 802.11AX CLIENTS

*Across Access Points and WLAN Switches Using IGMP Snooping
Audio/Video Multicast Use Cases and AV Bridge Interworking*

Open Compute Project, Time Appliances Project (TAP)

Version 0.11.1

Hesham Elbakoury, Independent

Ahmad Byagowi, Research Scientist, turbalance

Requirements & Strategy | White Paper | Shape Direction
Draft for Community Review | October 1, 2026

Date: October 1, 2026



Executive Summary

The Open Compute Project (OCP) Time Appliances Project brings together operators, implementers, silicon providers, and standards participants working on accurate and reliable distribution of time. This White Paper addresses a related multi-domain problem: maintaining low-disruption multicast delivery when IEEE 802.11ax clients roam across access points and WLAN switches, including professional audio/video systems that depend on IEEE 802.1 AVB/TSN timing and transport.

The paper combines event-driven IGMP membership repair, optional IEEE 802.11 group-delivery mechanisms, architecture-specific AVB/TSN boundary functions, IEEE 802.1AS time transfer, per-BSS TWT scheduling, and IEEE 1722 presentation timing. It distinguishes standardized building blocks from controller, endpoint, and boundary-function integration. It makes no universal guarantee of roam outage, packet loss, wireless latency, time error, or physical output simultaneity. Those properties require an explicit deployment budget and measurement.

Under the June 2026 OCP Contribution Taxonomy, the appropriate contribution type is Requirements & Strategy, White Paper, with the purpose of shaping direction. The intended audience is data-center and campus network operators, WLAN and switching vendors, professional AV engineers, and timing implementers. The paper is not an OCP Reference Architecture, Tested Configuration, Specification, or Benchmark because it does not yet include multiple validated configurations or a reproducible measurement dataset. Its value is a vendor-neutral architecture, a shared vocabulary, and a concrete set of validation metrics that TAP and Networking participants can refine collaboratively.

Date: October 1, 2026



This work is licensed under a [Creative Commons Attribution-ShareAlike 4.0 International License](https://creativecommons.org/licenses/by-sa/4.0/).

Table of Contents

1. Introduction.....	4
2. Compliance with Open Compute Project Tenets.....	4
3. Background.....	5
4. Candidate Architecture and Reference Topology.....	7
5. Mechanisms for Fast Multicast Re-Establishment.....	9
6. IEEE 802.11 Mechanisms for Wireless Multicast Reliability and Scheduling.....	11
7. Use Cases.....	12
8. Comparison of Wireless Multicast Delivery Mechanisms.....	16
9. Deployment and Validation Recommendations.....	16
10. Network-Wide Time Synchronization and Synchronized AV Delivery.....	17
11. Open Issues and Security Considerations.....	26
12. Future Work and OCP Contribution Evolution.....	28
13. Conclusion.....	29
14. Glossary.....	30
15. References.....	31
16. License.....	33
17. About Open Compute Foundation.....	34

1. Introduction

The Open Compute Project (OCP) Time Appliances Project (TAP) provides a forum for open, interoperable time distribution and time-sensitive applications. Multicast is the natural transport for one-to-many media distribution: a single stream reaches an arbitrary number of receivers without the sender replicating traffic per client. On switched Ethernet, IGMP snooping (RFC 4541) lets a Layer 2 switch forward each multicast group only to ports with interested receivers. IEEE 802.11, however, adds mobility, contention, group-key, and last-hop delivery behavior that complicate continuity when a receiver moves between access points.

This paper addresses a recurring deployment question: when 802.11ax clients roam among APs that may connect to different WLAN switches, how can multicast forwarding continuity and receiver presentation objectives be maintained while IGMP membership signaling and snooping drive wired forwarding state? The paper answers this at the architecture level (Sections 4 and 5), at the IEEE 802.11 mechanism level (Section 6), and through worked use cases that include professional AV bridging (Section 7).

2. Compliance with Open Compute Project Tenets

This contribution addresses all five OCP tenets as a best-effort White Paper. Statements below are design objectives and community value propositions, not measured product claims.

2.1. Openness

The architecture is based on published IEEE and IETF standards and separates standardized protocol functions from implementation-specific controller, boundary, and endpoint integration. No proprietary product is required. The paper is intended for multi-vendor review, extension, and validation across TAP and OCP Networking participants.

2.2. Efficiency

The design seeks to retain multicast's one-to-many efficiency while avoiding blanket conversion to per-receiver unicast. Event-driven membership repair can reduce stale or missing forwarding state, and selected reliability or scheduling mechanisms can be applied only where their measured benefit justifies airtime and control overhead. Any latency, airtime, or energy improvement remains deployment-specific.

2.3. Impact

The paper identifies a technical gap spanning WLAN mobility, multicast forwarding, and time-aware AV delivery. A common architecture and measurement vocabulary can help operators, equipment vendors, silicon providers, and standards participants compare implementations without converting one vendor's behavior into an assumed industry guarantee.

2.4. Scale

The candidate architecture covers multi-AP, multi-switch, and controller-managed domains. Its scale is conditioned on per-client and source-filter state, control-plane capacity, receiver count, retry

Date: October 1, 2026



This work is licensed under a [Creative Commons Attribution-ShareAlike 4.0 International License](https://creativecommons.org/licenses/by-sa/4.0/).

behavior, and time-reacquisition performance. The paper recommends explicit scalability and failure-mode evidence instead of fixed capacity claims.

2.5. Sustainability

Native multicast can reduce duplicate traffic compared with independent unicast copies when receivers share a stream, and event-driven membership repair may reduce unnecessary forwarding toward empty attachments. Energy effects also depend on transmit rate, retries, buffering, receiver count, and implementation. The paper therefore treats sustainability as a hypothesis to validate and recommends reporting airtime and energy alongside continuity and latency metrics.

3. Background

3.1. IGMP/MLD Snooping Fundamentals

IGMP snooping switches passively inspect IGMPv2/v3 (or MLDv1/v2 for IPv6) membership reports and queries between hosts and an upstream multicast router or querier, and use them to populate a per-VLAN table that maps a multicast group to listener-facing and multicast-router ports [7]. Frames addressed to a group are then forwarded only to those ports. IGMPv3 uses a default Query Interval (QI) of 125 s and Query Response Interval (QRI) of 10 s. RFC 9776 defines the Group Membership Interval as $\text{Robustness Variable} \times \text{QI} + 2 \times \text{QRI}$, which is 270 s at the default Robustness Variable of 2 [8]. An explicit state-change report may shorten removal, but fast-leave behavior is safe only when the implementation knows that no other listener remains on the port [20].

These defaults are tuned for a stable wired edge. Applied naively to a wireless edge, they create two failure modes during a roam: (1) at the new AP or port, no forwarding state exists until the next General Query or an explicit membership report is triggered, so the client sees a gap in the stream; and (2) at the old AP or port, stale forwarding state persists until timeout, wasting airtime and bandwidth if the switch keeps sending multicast toward an AP with no remaining listener for that group. Section 5 describes how proxy reporting and membership caching address failure mode (1), and how roam-triggered cleanup addresses failure mode (2).

3.2. 802.11ax Roaming-Relevant Amendments

IEEE 802.11ax was published as IEEE Std 802.11ax-2021 and is incorporated into IEEE Std 802.11-2024 [1], [2]. It improves spectral efficiency and spatial reuse through features such as OFDMA and BSS coloring, and it adds scheduled power-save operation through TWT. These features do not make EDCA channel access deterministic, move IGMP membership state between attachment points, or provide a universal roaming-latency bound. Any latency objective therefore remains conditional on scheduling, contention, retry behavior, and measured deployment performance.

- 802.11r (Fast BSS Transition, FT): reduces authentication and key-management signaling during a BSS transition and can shorten the handoff [1], [5]. The standard does not specify a universal sub-50 ms bound. Actual interruption also includes scanning, reassociation, data-

path updates, group-key delivery, and multicast-state repair, all of which must be measured for the deployment.

- 802.11k (Radio Resource Management): provides neighbor-report information that can help a client identify candidate APs and reduce scanning time before a roam [1], [6].
- 802.11v (BSS Transition Management and Directed Multicast Service, DMS): lets the infrastructure steer a client toward a target AP and lets a client request that selected group-addressed traffic be delivered as individually addressed frames [1], [4]. DMS can use unicast rate adaptation, acknowledgment, and retry behavior, but it does not eliminate loss and its airtime cost grows with receiver count (Section 6.2).

IEEE 802.11aa adds Groupcast with Retries (GCR), which provides feedback and retry options for group-addressed traffic (Section 6.1) [1], [3]. GCR Block Ack and GCR Unsolicited Retry can improve delivery when supported and configured, but neither makes group delivery equivalent to unicast nor supplies a universal loss or delay guarantee. Fast transition, membership repair, and last-hop reliability remain separate parts of the continuity design.

3.3. Problem Statement

Two 802.11ax clients, C1 and C2, are members of multicast group G, such as an IPTV channel or an AVB-mapped audio stream. At t_0 , both clients are associated with AP1 on WLAN Switch A, while AP2 on WLAN Switch B has no listener for G. C1 then roams from AP1 to AP2 while C2 remains associated with AP1. The design objective is to install membership and forwarding state for C1 at the new attachment point within a stated, measured outage and loss budget, while preserving delivery through AP1 for C2. Cleanup must therefore remove only C1's old per-station state and must not prune an aggregated AP uplink that still serves C2. Figure 1 defines this canonical scenario; later sections explain same-switch, cross-switch, and Layer 3 variants.




Figure 1. Canonical multicast roaming scenario. Before the roam, C1 and C2 receive group G through AP1 and Switch A while Switch B has no listener. After C1 roams to AP2, event-driven membership repair installs forwarding through Switch B, while Switch A retains forwarding for C2.

4. Candidate Architecture and Reference Topology

Figure 1 defines the cross-switch state transition used throughout this paper. The multicast source and querier sit above the distribution layer, WLAN Switches A and B maintain independent IGMP snooping state, and C1 moves from AP1 to AP2 while C2 remains behind AP1. The scenario therefore exercises both target-state installation at Switch B and safe retention at Switch A. Sections 4.1 through 4.3 apply the same state transition to same-switch, cross-switch, and Layer 3 anchored deployments.

4.1. Single-Switch, Multi-AP Domain

When both APs terminate on the same access switch and the same VLAN, that switch maintains the receiver-facing state for both AP ports. A roam requires it to add the new AP port to the group and to remove only the departing station's old state after confirming that no other listener remains behind the old AP uplink. Because both ports are local to one switch, this is the simplest case and needs no inter-switch multicast-routing change; event-driven proxy reporting can trigger re-registration faster than periodic query discovery, as discussed in Section 5.2.

4.2. Cross-Switch / Cross-Controller Domain

When AP1 and AP2 terminate on different access switches, each switch maintains independent snooping state. Continuity requires multicast reachability at both attachment points: either extend the receiver VLAN over a bridged trunk with an IGMP querier for that VLAN, or terminate separate receiver VLANs and route multicast between them, for example with PIM-SM and an explicitly designed rendezvous-point arrangement. In the routed case, each access switch directs its local membership toward a multicast-router port, and the distribution layer replicates G toward every switch with an active listener. The wireless-specific critical path is how quickly Switch B learns of C1 after the roam, which is where controller or AP proxy behavior matters most (Section 5.2).

In a centralized WLAN architecture, APs operating as Wireless Termination Points (WTPs) may communicate with an Access Controller (AC) using the CAPWAP base protocol in RFC 5415 and the IEEE 802.11 binding in RFC 5416 [18], [19]. CAPWAP supports Split MAC and Local MAC modes; it does not itself require a single data-plane termination point. When the relevant client traffic is centrally tunneled or bridged at the AC, IGMP state as seen by the wired network may remain on a stable controller uplink while the AC redirects the station's tunnel to the new AP. In a local-bridging design, membership remains an access-edge problem. In either case, whichever device is authoritative for group membership must learn of the roam and re-issue or redirect the join before traffic can resume.

4.3. Mobility Anchors for Layer 3 Roams

If AP1 and AP2 sit on different IP subnets, an architecture may preserve the client's original IP service through a mobility anchor. When that design anchors both multicast control and data, upstream membership state may remain stable while the tunnel endpoint changes. Other Layer 3 roaming designs must re-establish downstream membership and forwarding at the new attachment point. The deployment should therefore state where the client subnet, multicast control state, and data path are anchored rather than assuming that every Layer 3 roam is multicast-transparent.

5. Mechanisms for Fast Multicast Re-Establishment




Figure 2. Example event-driven multicast-state repair after a roam. The controller restores membership and forwarding state at the new AP; without an unsolicited or proxy report, query-driven discovery can approach QI + QRI.

5.1. Group Membership Caching

The access infrastructure should retain, for each stable station or controller-session identity, the multicast reception state learned from the client's IGMPv2/v3 reports. The cache should not blindly assume that a client MAC address is globally stable, because IEEE 802.11-2024 includes operation with randomized and changing MAC addresses [1]. The state model also needs to preserve IGMPv3 INCLUDE/EXCLUDE source filters rather than only a set of group addresses. Vendor implementations may expose related behavior as client tracking or multicast mobility.

5.2. Proxy-Join on Roam

When a client with cached group membership reassociates, the new attachment point (AP2 in Figure 2, or the controller acting on AP2's behalf) emits the appropriate IGMP membership report on the

new attachment path instead of waiting for the client to send another report or for the next scheduled General Query. With default timers, periodic discovery alone can leave a multi-second to multi-minute gap. This behavior is functionally similar to IGMP/MLD proxying (RFC 4605) at the access edge, but it is driven by a roam event rather than by a data-plane trigger and must preserve any IGMPv3 source-filter state.

5.3. Fast-Leave at the Old Attachment Point

The old AP/controller should recognize the client's departure and remove only that listener's state. The upstream switch port may be removed from a group only when no other listener remains behind it. RFC 9166 describes fast leave with the assumption that no more than one host is present on each VLAN port [20]. Because an AP uplink normally aggregates many stations, switch-port fast leave is unsafe unless the AP or controller supplies per-client tracking or equivalent listener state. This constraint prevents a departing station from pruning traffic still needed by another station on the same AP.

5.4. Timer Tuning

Where event-driven proxy reporting is unavailable, shorter query and membership timers can reduce stale-state duration and how long periodic discovery takes. They cannot ensure immediate discovery after association, because the next General Query may still be nearly a full Query Interval away. Table 1 lists the RFC 9776 defaults and example access-VLAN values that must be evaluated against control-plane load and receiver behavior [8].

Table 1. IGMPv3 timer defaults and example access-VLAN values for evaluation.

Parameter	IGMPv3 default	Example access-VLAN value	Rationale
Query Interval (QI)	125 s	30-60 s	Bounds the worst-case time for an unrefreshed listener to be re-discovered by a General Query
Query Response Interval (QRI)	10 s	2-3 s	Shortens the response window after a query; it does not replace QI in the new-listener discovery bound
Robustness Variable (QRV)	2	2 (unchanged)	Controls robustness against lost queries and reports; lowering it below 2 removes loss margin
Last Member Query Interval (LMQI)	1 s	0.25-0.5 s	Speeds leave verification when confirming that no listener remains on a port; it is not a join-discovery timer
Group Membership Interval (derived)	270 s	64-126 s for the example QI/QRI range	Router-side state timeout: $QRV \times QI + 2 \times QRI$; not a new-listener discovery bound

Aggressive timer tuning is a fallback, not a substitute for event-driven membership repair. For a newly attached listener that does not send an unsolicited report, the delay until a periodic query elicits a report can approach $QI + QRI$; the Last Member Query Interval applies to leave verification, not join discovery [8]. A controller-triggered proxy report can be much faster, but neither IGMP nor 802.11r supplies a universal end-to-end outage bound, so the result must be measured.

6. IEEE 802.11 Mechanisms for Wireless Multicast Reliability and Scheduling

Even with instantaneous wired-side snooping updates, wireless multicast remains sensitive to configured transmit rate, contention, interference, receiver diversity, and the selected feedback and retry mechanism. Legacy group-addressed delivery generally lacks per-receiver acknowledgment, but the rate is configurable and GCR or DMS changes the feedback and retry behavior. The following mechanisms therefore require deployment-specific configuration and measurement.

6.1. Groupcast with Retries (GCR) (IEEE 802.11aa)

GCR defines feedback and retry options for group-addressed delivery, including GCR Block Ack (GCR-BA) and GCR Unsolicited Retry (GCR-UR) [1], [3]. GCR-BA can use member feedback to retransmit frames, while GCR-UR repeats frames without soliciting per-member acknowledgments. Either mode can improve delivery when it is supported and configured, but neither creates a universal loss or latency guarantee: performance depends on receiver count, channel conditions, retry policy, and implementation. GCR membership and IP multicast membership are distinct state and must both be kept consistent as a station roams; the AP maintains the GCR roster and updates it on each roam-in/roam-out event, in addition to (not instead of) the IGMP group-membership cache described in Section 5.

6.2. Directed Multicast Service (DMS) (IEEE 802.11v)

DMS lets a client request that an AP deliver selected group-addressed traffic as individually addressed frames [1], [4]. Those copies can use unicast rate adaptation, acknowledgment, and retry behavior, which may improve delivery for a small receiver population. DMS does not eliminate loss, and its airtime cost grows with the number of receivers because the AP transmits a separate copy to each one. It is therefore a deployment choice for cases whose measured reliability benefit justifies that cost; typically BSSs with a small number of concurrent multicast receivers per AP, or latency/loss-critical use cases such as live conferencing or paging, discussed in Section 7.3.

6.3. Target Wake Time (TWT) and Multicast Scheduling

Broadcast TWT lets an AP advertise wake and service periods for groups of stations within its BSS [1]. This can structure delivery opportunities and reduce avoidable receiver wake time, but it does not reserve the channel or by itself bound packet transmission or arrival. EDCA contention, channel availability, interference, retries, and AP implementation remain in the delay budget. Cross-AP alignment, which is needed for two APs carrying the same stream to schedule near one another, requires additional controller coordination, as discussed in Section 10.2.

6.4. BSS Coloring and Spatial Reuse

802.11ax BSS coloring lets an AP mark its transmissions with a color that neighboring BSSs on the same channel can use to distinguish co-channel interference from an overlapping BSS, enabling more aggressive spatial reuse. This is not multicast-specific, but in dense multi-AP deployments carrying the same multicast stream on multiple adjacent BSSs (the common case for whole-building AV/paging systems), reduced co-channel contention directly improves the achievable GCR/DMS retry budget and reduces jitter, which is why it is included here as a supporting mechanism rather than the primary one.

7. Use Cases

Figure 3 provides a visual index to the four use cases. The multicast-control problem is shared across the panels, but the topology, delivery policy, and latency tolerance differ. The panels are illustrative deployment patterns, not standardized service guarantees.




Figure 3. Four use-case patterns. Panels A through D distinguish same-switch buffered IPTV, cross-switch streaming that requires new forwarding state, low-loss paging using DMS as an example last-hop policy, and professional AV with an architecture-specific boundary function.

7.1. Use Case A: Enterprise IPTV / Digital Signage (Single-Switch Domain)

As summarized in Figure 3, panel A, a retail or campus building distributes an IPTV channel or signage feed to displays associated with different APs on the same access switch and VLAN. Roaming is infrequent. This is the simplest case in Section 4.1: one switch maintains the relevant snooping state, membership is tied to a stable station or controller-session identity, and an optional wireless reliability mechanism such as GCR-BA may be used where supported. Buffered video may tolerate a larger measured repair interval than interactive media, but proxy reporting and last-hop delivery behavior still require validation.

7.2. Use Case B: Cross-Switch Roaming for Streaming Clients

Figure 3, panel B, summarizes a larger campus in which AP1 and AP2 terminate on different access switches, as described in Section 4.2. The receiver VLAN may be bridged across the distribution layer, or separate receiver VLANs may be connected by multicast routing. The deployment-specific requirement is to provision multicast reachability consistently at every candidate attachment point and to install the new switch's listener state promptly after a roam. A common failure is a VLAN or multicast-routing policy present on Switch A but absent at Switch B, which breaks continuity only after a cross-switch roam.

7.3. Use Case C: Real-Time Audio/Video Conferencing and Paging

Figure 3, panel C, summarizes live, two-way, or safety-critical audio, such as emergency paging, push-to-talk, or a conferencing multicast fan-out. These applications combine low latency and low loss tolerance and may have a small receiver set per BSS. DMS unicast conversion at the AP is one possible last-hop policy where supported. IGMP snooping still governs the wired replication decision, determining which APs receive the stream, while DMS governs how the selected AP delivers it over the wireless hop. Both the multicast repair interval and the last-hop reliability behavior require measurement.

7.4. Use Case D: Professional AV Bridging with IEEE 802.1 AVB/TSN

Figure 3, panel D, introduces the professional AV pattern developed in Figures 4 and 5. Professional AV systems may use IEEE 802.1 AVB/TSN profiles, including IEEE 802.1AS time, IEEE 802.1Q stream registration and shaping, and IEEE 1722 transport [11]-[13], [16]. An IEEE 802.11 access link does not inherently require a separate AVB gateway if the AP and stations implement the applicable AVB/TSN profiles. The architecture considered here instead treats an ordinary enterprise WLAN as a domain boundary. A boundary function may relay or terminate and recreate AVTP, reservation, and timing functions; translate multicast membership; and map traffic into a WLAN QoS policy. IEEE 1722-2025 supports native IEEE 1722 frames and IEEE 1722 IP/UDP encapsulated frames [16], but a specific AVTP-to-IP-multicast, SR-class-to-WMM, and admission mapping remains a profile or boundary-function design choice. Where such a function is used, it maps each AVTP stream or flow to the selected IP multicast group, maps its SR class separately to an architecture-specific WLAN policy, and implements the group-membership behavior described in Section 5.




Figure 4. Example boundary architecture for an AVB/TSN talker and an IGMP-snooped 802.11ax segment. The optional boundary function maps or recreates stream, membership, timing, and WLAN QoS policy; wired reservation and latency properties do not automatically extend across the WLAN.

A wired AVB reservation or latency target does not automatically extend across a conventional WLAN. WMM priority, GCR/DMS reliability, and TWT service-period structure can support a service objective, but they do not create an IEEE 802.1Q-equivalent wireless delay guarantee. IGMP snooping determines where the IP multicast stream is forwarded; it does not preserve AVB stream identity or reservation semantics by itself. Figure 5 and Table 2 therefore show an architecture-specific mapping whose performance must be validated.




Figure 5. Example, non-standardized mapping from AVB Stream Reservation classes to WMM access categories. The 2 ms and 50 ms values are wired AVB profile path targets only; GCR or DMS does not preserve those targets across the WLAN.

Table 2. Example architecture-specific mapping of AVB Stream Reservation classes to 802.11e/WMM access categories.

AVB SR Class	Wired profile target	802.1Qav priority	Example WLAN policy	Typical stream
Class A	2 ms across the profiled wired topology	Highest credit-based queue	AC_VO plus selected reliability; no corresponding WLAN guarantee	Live audio, in-ear monitor
Class B	50 ms across the profiled wired topology	Medium credit-based queue	AC_VI plus selected reliability; no corresponding WLAN guarantee	Camera feed, video program
Best-effort multicast	Not guaranteed	No shaper (standard IGMP VLAN)	AC_BE	Signage, non-real-time IPTV

8. Comparison of Wireless Multicast Delivery Mechanisms

Table 3 compares the mechanisms discussed in Section 6 across the dimensions most relevant to a roaming, multi-switch deployment: delivery reliability, airtime cost, sensitivity to roam events, and scalability with receiver count.

Table 3. Qualitative trade-offs among wireless multicast delivery mechanisms for roaming 802.11ax clients.

Mechanism	Delivery behavior	Airtime cost	Roam dependency	Receiver scaling
Native multicast plus IGMP snooping	No per-frame receiver feedback and no delivery guarantee	Lowest: one copy per AP	The new AP still needs wired forwarding and listener state	Excellent
GCR Unsolicited Retry (GCR-UR)	Fixed repeats without member feedback; improvement depends on the repeat and loss budget	Low to medium	The new AP must restore forwarding and applicable GCR state	Very good
GCR Block Ack (GCR-BA)	Member feedback and retries where supported; no universal loss or delay bound	Medium to high; grows with feedback and retries	The new AP must restore GCR membership and forwarding state	Depends on group size and implementation
DMS (individually addressed delivery)	Per-receiver ACK and retry behavior where supported; loss remains possible	High: one copy per receiver	DMS and forwarding state must remain valid or be restored after the roam	Limited at high receiver counts

9. Deployment and Validation Recommendations

- Enable IGMP snooping on the access switches, ensure that every receiver VLAN has an active IGMP querier, validate querier redundancy, and verify that the multicast VLAN or PIM rendezvous-point reachability is provisioned on every switch to which a client may roam, not only on the switches used in initial testing.
- Prefer controller-based or otherwise coordinated AP deployments when they can drive membership caching, proxy reporting, and per-listener cleanup directly from the 802.11 roam event. Do not rely solely on generic switch timer tuning.
- Use 802.11r, 802.11k, and 802.11v where supported to reduce mobility delay. Evaluate GCR-BA for buffered or non-interactive AV multicast and DMS for small-receiver-count, latency-sensitive streams such as live paging or conferencing. Treat both as optional mechanisms whose benefit and airtime cost must be measured.
- For the assumed non-AVB enterprise WLAN, use an explicit boundary function with documented AVTP, reservation, timing, IP-multicast, and WLAN QoS behavior. Whether the function relays or terminates and recreates gPTP and AVTP is architecture-specific. Do not describe it as preserving wired AVB bounds unless an end-to-end profile and measurements establish those bounds.
- Broadcast TWT may structure multicast wake and service periods within one BSS, but it does not reserve the channel or by itself bound jitter. Cross-AP alignment requires controller

Date: October 1, 2026



This work is licensed under a [Creative Commons Attribution-ShareAlike 4.0 International License](https://creativecommons.org/licenses/by-sa/4.0/).

coordination and a maintained gPTP-to-TSF mapping at each AP. Contention, interference, retry delay, and scheduler error remain terms in the validated budget.

- Tune IGMP snooping timers using Table 1 only as a fallback for deployments without event-driven repair. Measure the resulting discovery delay and control-plane load; do not treat timer values alone as an end-to-end roaming bound.

For OCP review, each implementation should report at least the measured roam-to-data outage, loss behavior, maximum time error during the roam, recovery time, end-to-end latency budget, and residual presentation skew. Results should identify receiver count, channel load, reliability mode, TWT policy, security mode, and topology so other implementers can reproduce the conditions.

10. Network-Wide Time Synchronization and Synchronized AV Delivery

Sections 4 and 5 describe how multicast state can be repaired after a roam, and Section 6 discusses wireless delivery mechanisms. Those mechanisms do not establish presentation synchronization or an end-to-end AV latency guarantee. This section identifies the standardized capabilities and additional implementation assumptions needed to distribute a common time base, target coordinated presentation, and construct a testable latency budget. It does not claim uninterrupted synchronization or a standards-backed wireless latency guarantee.

10.1. Network-Wide Time Synchronization Across APs and Switches

Within an infrastructure BSS, the IEEE 802.11 Timing Synchronization Function (TSF) supplies the local MAC time base used for beacon scheduling and TWT. TWT target times are expressed relative to the serving BSS's TSF timer. TSF counter resolution and BSS synchronization do not establish either a cross-BSS time base or a guaranteed relationship to IEEE 802.1AS time [1].

Independent APs are not required to maintain phase-aligned TSF timers. TSF is therefore useful for BSS-local scheduling, but TSF alone is not evidence of network-wide media synchronization. Any claimed offset or jitter bound must come from a specified synchronization method and measured conformance, not from the one-microsecond TSF counter unit.

IEEE Std 802.1AS-2025 defines transport of synchronized time over supported LAN media [12]. For IEEE 802.11 links, its WLAN media-dependent procedures can use IEEE 802.11 Timing Measurement (TM) or Fine Timing Measurement (FTM), subject to peer capability and implementation support [15]. These procedures use MAC/PHY timestamps and upstream synchronization information to support link-delay measurement and time transfer. Ordinary FTM ranging support does not, by itself, make an AP or STA an 802.1AS time-aware system or establish a synchronization-error bound. Timestamp placement, delay asymmetry, measurement noise, cadence, filtering, and clock control all remain part of the implementation and validation.




Figure 6. IEEE 802.1AS use of IEEE 802.11 FTM. The timeReceiver uses transmitter and local timestamps, together with upstream synchronization information, to derive synchronized time and mean link delay. Achieved error and convergence time are implementation- and channel-dependent.

For each exchange, t_1 and t_4 are captured in the timeTransmitter time base and t_2 and t_3 in the timeReceiver time base. The timeTransmitter conveys t_1 and t_4 , together with synchronization information associated with the preceding measurement, to the timeReceiver. On an IEEE 802.11 link, the timeReceiver uses the link-delay result; multiple measurements may be filtered before the required accuracy is reached [12], [15]. The four timestamps alone do not establish a universal sub-microsecond or per-roam convergence claim.

Candidate architecture: explicit support and configuration are required at every element




Figure 7. Candidate timing architecture. A time-aware wired network supplies a common gPTP time base to capable APs; each STA requires its own IEEE 802.1AS PortSync and IEEE 802.11 media-dependent instance. Every element shown must explicitly support and be configured for this function.

If the selected time source, intervening time-aware systems, APs, and STAs all implement the required IEEE 802.1AS functions and are in steady state, they can maintain synchronized local clocks on a common gPTP timescale [12]. This does not necessarily imply UTC traceability. The achievable maximum time error depends on topology, timestamp quality, asymmetry, filtering, and endpoint clock control, and must be validated for the deployment.

Roaming does not automatically carry timing continuity the way it carries multicast group membership in Section 5. The target AP may already be synchronized, but IEEE 802.11r does not preserve IEEE 802.1AS PortSync state or require TM/FTM reacquisition; a roam creates a new association and a new timing relationship [1], [5], [12]. A controller may trigger timing reacquisition after reassociation (for example, alongside the 802.11r fast-transition handshake), but continuity then depends on STA holdover, AP-to-AP time error, measurement and servo convergence, and handling of phase or frequency discontinuities. The design should treat this as a measured transient rather than an assumption: it needs an explicit holdover, reacquisition, and failure policy, with the maximum time error during the roam and the time to return to the required bound measured and documented, not inferred from the steady-state behavior described above.

Date: October 1, 2026



This work is licensed under a [Creative Commons Attribution-ShareAlike 4.0 International License](https://creativecommons.org/licenses/by-sa/4.0/).

PTP Role Assignment Across the WLAN

IEEE Std 1588-2019 defines PTP clock types, including ordinary, boundary, and transparent clocks; IEEE Std 802.1AS-2025 defines the gPTP profile and its time-aware-system model [12], [21]. Current 802.1AS terminology includes timeTransmitter and timeReceiver, PTP End Instance, and PTP Relay Instance. Mapping a switch or AP to one of these functions is a deployment assumption, not an inherent property of the device.

Table 4. IEEE 802.1AS functions in the candidate WLAN timing topology.

IEEE 802.1AS function	Candidate element	Status and limit
Selected time source / best timeTransmitter	Dedicated source or capable core/distribution element	Selected or configured for the gPTP domain; GNSS and UTC traceability are optional
PTP Relay Instance	Only switches and APs that explicitly implement the required time-aware relay and media-dependent functions	Do not label every switch or AP a boundary clock by assumption
PTP End Instance / timeReceiver	STA with PortSync and IEEE 802.11 media-dependent support for its association	TM or FTM capability plus an IEEE 802.1AS implementation is required
Transparent clock	Not an IEEE 802.1AS conformance role	Interworking with a non-gPTP IEEE 1588 segment needs an explicit boundary design

IEEE 802.1AS uses PTP Relay Instances rather than generic IEEE 1588 transparent-clock operation. A relay instance conforms to the IEEE 1588 boundary-clock model, although its synchronized transport may be mathematically equivalent to a peer-to-peer transparent clock [12], [21]. Calling every bridge a boundary clock is therefore a useful analogy, not a complete conformance statement. Interworking with a non-gPTP IEEE 1588 segment is not automatic and requires an explicit profile/domain-boundary design.

PTP/gPTP vs. NTP

NTP may coexist with gPTP, but it does not implement the IEEE 802.1AS profile and cannot substitute for the time base used by this paper's IEEE 1722 presentation-time design. Table 5 compares protocol roles without assigning universal accuracy numbers, because performance depends on profile, timestamping, topology, implementation, and operating conditions.

Table 5. IEEE 802.1AS/gPTP and NTP roles in the proposed AV media-timing architecture.

Dimension	IEEE 802.1AS / gPTP	NTP
Protocol relationship	IEEE 1588 profile for time-sensitive IEEE 802 networks	IP time-synchronization protocol
IEEE 802.11 operation	Media-dependent procedures can use TM or FTM when both peers support them	Runs as IP traffic; NTPv4 defines no IEEE 802.11 media-dependent timestamp transport
IEEE 1722 relationship	IEEE 1722 presentation timing uses the applicable gPTP timescale	IEEE 1722 defines no NTP presentation-time mapping
Performance statement	Maximum time error depends on profile, topology, timestamps, and implementation	Error depends on servers, network path, timestamping, and implementation

Date: October 1, 2026



This work is licensed under a [Creative Commons Attribution-ShareAlike 4.0 International License](https://creativecommons.org/licenses/by-sa/4.0/).

Dimension	IEEE 802.1AS / gPTP	NTP
Use in this architecture	Media time after conformance and performance validation	Management time or loose-tolerance application time

The protocols can coexist: NTPv4 may serve management and general infrastructure timekeeping [24], while IEEE 802.1AS supplies the gPTP timescale used by an IEEE 1722 media design [12], [16]. Whether NTP is sufficient for loose-tolerance IPTV or signage is an application decision based on a stated skew budget. It is not a conclusion that follows from AVB Class A or Class B labels.

10.2. Synchronized Multicast Delivery Regardless of Client Location

With a validated common gPTP timescale, presentation timestamps can provide a shared receiver timing reference. Coordinated radio scheduling is optional and does not itself guarantee simultaneous delivery or presentation. The design therefore separates per-BSS scheduling from receiver presentation timing.

Coordinated scheduling via broadcast TWT

Broadcast TWT is negotiated and expressed in the serving BSS's local TSF time [1]. It identifies wake and service periods; it does not reserve the channel, guarantee a transmission opportunity, or define coordination among independent APs. A controller could select a gPTP target T_g , convert T_g into each AP's local TSF using a maintained mapping, and arrange overlapping service periods for a multicast group. That is a proposed implementation, not a cross-AP TWT procedure standardized by IEEE 802.11. Its error budget includes gPTP error, mapping error, TWT scheduling granularity, medium-access delay, and retries. The active IEEE P802.1Qee project is relevant to traffic engineering across wireless links with greater delay variance, but it is not yet a published standard [23]. Figure 8 illustrates overlapping service periods mapped from a common gPTP target, which may narrow but cannot guarantee the spread in radio transmission time.




Figure 8. Proposed, non-standardized cross-AP scheduling. A controller maps common gPTP targets into each BSS's TSF and arranges overlapping broadcast-TWT service periods. Presentation timestamps, rather than simultaneous radio transmission, provide the receiver timing reference.

Common presentation time at the receiver

For an IEEE 1722 stream format that carries an AVTP timestamp, the talker associates the timestamp with a designated media sample; the listener maps it to its local gPTP timescale and presents the media according to the stream and profile rules [12], [16]. Data that arrives early can be buffered until the target. The timestamp does not guarantee identical audible or visible output: clock error, timestamp interpretation, listener scheduling, codec or resampling delay, and DAC/display pipelines require a separate bound or calibration.

This architecture can support bounded presentation skew only if multicast forwarding, wireless delivery, time synchronization, timestamp interpretation, listener buffering, and endpoint processing all meet validated limits. IGMP snooping does not guarantee packet delivery, TWT does not guarantee simultaneous transmission, and an AVTP timestamp alone does not guarantee acoustic or display simultaneity.

Buffer Release Synchronization Mechanism

A presentation timestamp identifies a target at the media-application interface. A listener buffer and release scheduler turn that target into a presentation action, while a separate media-clock-recovery or Clock Reference Format function controls long-term rate. Actual output alignment also includes downstream processing and calibration. These mechanisms should be budgeted separately.

Date: October 1, 2026



This work is licensed under a [Creative Commons Attribution-ShareAlike 4.0 International License](https://creativecommons.org/licenses/by-sa/4.0/).

Per-frame release comparator

For a timestamped stream, each receiver maps the designated sample time to its synchronized local gPTP clock and schedules release at the media-application interface. Receivers target the same presentation time, but the actual inter-listener output skew is bounded by combined gPTP error, timestamp granularity, listener scheduling, and calibrated downstream processing. It is therefore inaccurate to say that every comparator fires at exactly the same physical instant.




Figure 9. Example listener timing policy. Each receiver targets the same presentation time for a designated sample; actual output skew remains implementation- and calibration-dependent. Panel B shows one possible late-frame concealment policy, not a requirement of IEEE 1722.

Presentation offset and receiver buffer depth are related but not identical. The offset is selected so that data is expected to reach the listener before presentation; receiver buffering also includes local pipeline and processing constraints. Both require a validated worst-case delay, or an explicitly chosen statistical objective, plus margin. IEEE 1722 does not supply a universal buffer size, and buffering does not make an otherwise unbounded wireless path deterministic.

Media clock recovery

IEEE 1722 supplies presentation-timing semantics and, where selected, Clock Reference Format streams [16]. The design of a PLL, resampler, or other media-clock-recovery loop is implementation- and profile-specific. Its acquisition time, wander rejection, and response to a roam or time

Date: October 1, 2026



This work is licensed under a [Creative Commons Attribution-ShareAlike 4.0 International License](https://creativecommons.org/licenses/by-sa/4.0/).

discontinuity contribute to the endpoint error budget; IEEE 1722.1 should not be cited as specifying that control loop.

Handling a missed deadline

The response to a late AVTPDU is endpoint- and application-specific. Dropping, concealing, muting, interpolating, or controlled resynchronization are possible policies, but IEEE 1722 does not mandate the policy illustrated here. Figure 9 Panel B should therefore be read as an example that preserves a chosen presentation schedule, not as a universal receiver rule.

Table 6 distinguishes standardized presentation semantics from implementation-specific receiver behavior.

Table 6. Presentation and receiver mechanisms, with their standardization status and design purpose.

Mechanism and status	What it controls	Limit if omitted or unbounded
AVTP presentation timestamp (standardized semantics)	Target presentation of a designated sample at the media-application interface	No shared presentation target
Listener buffer and release scheduler (implementation)	Absorbs early-arrival variation until the target	Late data or excessive added latency
Media-clock recovery or CRF processing (format/profile plus implementation)	Long-term sample-rate relationship	Buffer drift, resampling error, or loss of lock
Output-delay calibration and late-frame policy (implementation)	Actual output alignment and response to a missed deadline	Unbounded output skew or inconsistent recovery

10.3. Maintaining End-to-End AV Latency Requirements

The familiar 2 ms Class A and 50 ms Class B values are AVB bridged-LAN profile latency targets for a bounded wired topology, not per-hop values and not a complete talker-to-loudspeaker or camera-to-display budget [11], [13]. An IEEE 1722 presentation offset is configurable and is not universally equal to either value [16]. Introducing an IEEE 802.11 segment does not automatically extend the wired AVB latency contract (Table 2). Meeting a chosen end-to-end target once a wireless hop and possible roaming are in the path requires accounting for every source of delay between talker and listener and bounding each one explicitly, rather than assuming the wired profile's target simply carries through the AP.

For a mixed wired/wireless path, Table 7 is an engineering budget rather than a standards guarantee. Every term needs a defensible analytical or measured worst case, and the required presentation margin must cover the selected service objective. WMM priority and TWT placement alone do not establish a deterministic wireless delay bound.

Table 7. Engineering latency budget for a mixed wired/wireless AV path.

Path term	Basis for the budget	Required evidence
Talker encode / packetize	Endpoint implementation	Measured worst case for the selected codec and packetization

Path term	Basis for the budget	Required evidence
Wired AVB path	Applicable IEEE 802.1Q shaping and AVB profile within a conformant topology	Use a standards-backed bound only inside that profiled wired domain; MSRP registers and admits the stream, while shaping and profile rules supply the delay analysis
Boundary-function processing (if present)	Architecture-specific relay or translation implementation	Measured relay, translation, queueing, and timestamp-handling delay
Wireless queue and medium access	WMM/TWT priority and service-period structure	No standardized worst-case WLAN delay; analyze and measure queueing, contention, scheduler error, and interference
Wireless delivery and retries	Configured GCR/DMS mode and retry policy	Measure loss, retry delay, and receiver-scaling effects
Roam outage and timing reacquisition	FT plus data-path repair, group-key readiness, multicast-state repair, and separate TM/FTM reacquisition	No universal FT or timing-convergence bound; measure worst-case outage, maximum time error, and return-to-bound time
Listener buffer and endpoint output	Selected presentation margin plus calibrated processing delay	Absorbs only bounded delay; report residual presentation skew and miss policy

MSRP registers stream attributes and supports admission within the IEEE 802.1Q bridged domain; the applicable shaping and AVB profile provide the wired delay analysis [13]. Neither function standardizes admission or a delay bound for an IEEE 802.11 hop. A boundary function may implement wireless-aware admission using measured queueing, access, retry, and roam outage, but that extension is implementation-specific. IEEE P802.1Qee is an active project addressing traffic engineering for bridged networks that include wireless technologies with greater delay variance; it does not yet provide a published guarantee [23]. A possible tens-of-milliseconds roam cannot fit within a 2 ms Class A network target. Reclassification must be configured end to end; alternatives include greater application buffering, redundant or make-before-break delivery, or an explicitly accepted interruption.

10.4. Standardized Building Blocks vs. Required Custom Engineering

Sections 10.1 through 10.3 assemble capabilities from IEEE 802.1AS, IEEE 1722, IEEE 802.11, and IGMP/MLD, but their status differs: some are standardized protocol functions, some are optional device capabilities, and others are proposed controller, boundary-function, or endpoint integration. Table 8 identifies those limits so that standards compliance is not mistaken for a complete synchronized WLAN service.

Table 8. Standards-based capabilities used by the architecture and their limits.

Capability	Standard basis, status, and limit
Common gPTP timescale	IEEE 802.1AS-2025 plus a supported IEEE 802.11 TM or FTM media-dependent procedure; ordinary FTM capability alone does not establish the service or maximum time error
AVTP presentation timing	IEEE 1722-2025 for applicable stream formats; endpoint pipeline remains implementation-specific
Fast transition and radio management	IEEE 802.11-2024 FT/RRM/BTM; no universal roam-latency or timing-reacquisition bound

Date: October 1, 2026



This work is licensed under a [Creative Commons Attribution-ShareAlike 4.0 International License](https://creativecommons.org/licenses/by-sa/4.0/).

Capability	Standard basis, status, and limit
Wireless multicast delivery	IEEE 802.11-2024 GCR/DMS/TWT as applicable; optional capabilities do not supply an end-to-end latency guarantee
Roam-aware multicast state	IGMP/MLD snooping and proxying plus vendor/controller event integration; fast leave needs per-listener state

No single published standard coordinates all of these functions across APs. Cross-AP TWT mapping, wireless-aware admission, roam-triggered membership and timing actions, and endpoint output calibration remain integration work. IEEE P802.1Qee directly addresses part of the wireless traffic-engineering gap, but as an active project with no published standard it should be described as relevant ongoing work, not available functionality [23]. Table 9 identifies the remaining engineering boundaries.

Table 9. Integration gaps and the component responsible for closing each one.

Integration gap	Why the standards do not close it	Engineering owner
Cross-AP gPTP-to-TSF/TWT mapping	IEEE 802.11 defines TWT per BSS; it does not define cross-AP scheduling or a channel-access bound	WLAN controller/AP scheduler plus analytical and measured validation
Wireless-aware admission and delay bound	MSRP registers and admits streams in the bridged domain; applicable credit-based shaping and the AVB profile supply wired delay analysis. Neither models WLAN contention, retry, TWT scheduling, or roam outage	Boundary function and WLAN controller using measured or analytical limits for the selected products and topology
Endpoint presentation and late-frame behavior	IEEE 1722 supplies timestamp semantics, not downstream output calibration or a universal concealment policy	Endpoint media pipeline and system calibration
Roam-triggered timing and membership repair	FT does not require TM/FTM reacquisition or a proxy IGMP report	AP/controller integration, with holdover and failure policy
Endpoint output calibration	IEEE 1722 specifies presentation timing at the media interface, not the downstream DAC or display path; actual output skew is device-specific	Endpoint hardware or firmware calibration, validated per device

The practical implication is that this architecture requires capability discovery, configuration, and measurement, not merely feature enablement. A WLAN controller may coordinate per-AP schedules and roam events; an AV boundary function may implement stream and admission policy; and endpoints must implement and calibrate presentation behavior. The resulting maximum time error, packet-delay distribution, roam outage, and presentation skew should be stated as measured system properties.

11. Open Issues and Security Considerations

The architecture in Sections 4 through 10 depends on mechanisms whose security properties, implementation support, and integration vary by deployment. Three issues require explicit treatment: multicast group-key availability during a roam, trust and integrity of the selected TM/FTM and IEEE 802.1AS timing procedures, and the difference between standards-defined capability and product support.

Date: October 1, 2026



This work is licensed under a [Creative Commons Attribution-ShareAlike 4.0 International License](https://creativecommons.org/licenses/by-sa/4.0/).

11.1. Multicast Key Management Across a Roam

Sections 4 and 5 address multicast continuity at the forwarding layer by ensuring that the right frames reach the right AP, but they do not address decryption. Group-addressed IEEE 802.11 frames are encrypted with a Group Temporal Key (GTK), delivered to each station through the 4-way handshake or the dedicated Group Key Handshake during rekey, and normally scoped per BSS. A client that roams to AP2 cannot decrypt multicast or broadcast traffic there until it holds AP2's GTK. This requirement is distinct from holding valid unicast keys or completing IEEE 802.11r Fast BSS Transition authentication.

Fast BSS Transition reassociation can carry GTK and IGTK key data in the Reassociation Response when the selected security mode and implementation support that behavior [1], [5]. This is not a universal FT guarantee. A deployment that requires multicast continuity should verify that group keys are delivered and installed before multicast resumes; otherwise a post-roam decryption gap can remain even after forwarding state has been repaired.

GTK rekeying is independent of roaming but can create loss if a station has not installed the new key before cutover. The deployment should measure rekey behavior for active multicast streams and coordinate key transitions across the mobility domain where supported. TWT can help schedule receiver wake opportunities, but placing rekey outside a service period does not guarantee a loss-free transition and is not required by IEEE 802.11.

11.2. FTM/Ranging Trust and Integrity

FTM frame formats and procedures are specified in IEEE Std 802.11-2024, Clauses 9.6.7.32-9.6.7.33 and 11.21.6 [15]; protected ranging extensions are defined separately in IEEE 802.11az-2022 [17]. The IEEE 802.1AS WLAN media-dependent procedure is specified in IEEE Std 802.1AS-2025, Clause 12 [12]. A positioning exchange should not be assumed to provide every integrity property needed for time transfer, but its security properties should not be generalized to every IEEE 802.1AS implementation. The deployment must identify the exact procedure, peer-authentication assumptions, and management-frame-protection assumptions, then validate behavior under spoofed, replayed, or manipulated timing inputs.

IEEE 802.11az-2022, incorporated into IEEE 802.11-2024, defines protected ranging features [17]. Their applicability to the IEEE 802.1AS WLAN media-dependent procedure, product support, and interoperability must be verified. Where timing integrity matters, select a documented protected procedure and include peer authentication, source selection, anomaly detection, holdover, and failure behavior in the validation plan.

11.3. Standards vs. Shipping Silicon: A Practical Caveat

Every mechanism in Sections 6 and 10 is described at the level of what the relevant IEEE standard permits. Actual availability in shipping AP/controller silicon and software is uneven: GCR-BA membership-management interfaces, cross-AP broadcast-TWT coordination hooks, and IEEE 802.1AS WLAN synchronization support are not uniform across vendors or product generations. A deployment built from this architecture should be validated against the selected platform through vendor documentation, interoperability testing, or direct engagement with the vendor's roadmap.

Date: October 1, 2026



This work is licensed under a [Creative Commons Attribution-ShareAlike 4.0 International License](https://creativecommons.org/licenses/by-sa/4.0/).

Standards compliance establishes what is possible; it does not establish what a particular product exposes to an integrator.

12. Future Work and OCP Contribution Evolution

The following extensions were left out of the scope of this paper but would materially strengthen it as a deployment reference rather than an architectural description:

- Quantitative capacity and scalability modeling: derive closed-form or simulation-based bounds on maximum GCR-BA membership per AP before retry overhead saturates airtime, and size TWT service periods as an explicit function of receiver count and retry budget instead of relying on the qualitative comparisons in Table 3.
- IPv6 and MLD parity: extend the paper's IPv4-focused analysis to MLDv1/v2 snooping. MLD is the IPv6 counterpart to IGMP, but dual-stack and IPv6-only AV-over-IP deployments need explicit treatment of membership state, proxy reporting, and timer behavior.
- Fault-tolerant timing and time integrity: IEEE 802.1ASed-2026 standardizes mechanisms that use two or more gPTP domains, multiple time-distribution paths, local-oscillator input, and a time-selection function [22]. Remaining work for this architecture includes defining domain and path independence, selection thresholds, holdover behavior, and maximum transient time error for the WLAN topology, including during a station roam, and coordinating that plan with IGMP querier redundancy (Section 9).
- Wireless mesh and repeater backhaul: analyze topologies that introduce additional multicast-forwarding and timing segments. Each segment needs an explicit forwarding authority and, where synchronized media is required, a supported time-aware function with a measured contribution to maximum time error.
- IEEE 802.11be-2024 (Wi-Fi 7) Multi-Link Operation: evaluate its use for redundant multicast continuity and its interaction with GCR, DMS, timing reacquisition, and duplicate elimination. The standard is published [25], but those system-level behaviors remain future work for this paper.
- 6 GHz-specific reliability: validate multicast rate selection, retry assumptions, regulatory power constraints, and channel conditions for the selected 6 GHz deployment. High Efficiency operation is required in the 6 GHz band, but that requirement does not remove the need for deployment-specific reliability measurements [1].

12.1. Follow-on OCP Contributions

This White Paper can evolve into an OCP Reference Architecture only after at least two distinct configurations are validated with complete configuration, test, and certification evidence [26]. A later Benchmark could publish collected metrics and a reproducible method. Controller, boundary-function, or measurement implementations could be contributed separately as Source Code. These are potential future artifacts, not claims made by the current contribution.

Date: October 1, 2026



This work is licensed under a [Creative Commons Attribution-ShareAlike 4.0 International License](https://creativecommons.org/licenses/by-sa/4.0/).

13. Conclusion

Low-disruption multicast delivery to roaming 802.11 clients can be supported when the roam event actively drives group-membership repair through per-client state and a proxy report at the new attachment point. Cleanup at the old point must preserve traffic for any other listener behind the same AP. GCR and DMS can improve last-hop reliability when supported, but continuity remains a measurable service property rather than a guarantee supplied by IGMP snooping or fast transition alone.

For professional AV, a common gPTP timescale is conditional on conformant participating PTP Instances, steady-state synchronization, and a validated maximum time error [12]. IEEE 802.11 TM/FTM can support the 802.1AS WLAN media-dependent procedure, but generic ranging capability does not establish that service. Broadcast TWT supplies per-BSS wake/service timing, not a standardized cross-AP channel-access bound. IEEE 1722 presentation timing can align a designated sample at the media-application interface when packets arrive before their deadlines; endpoint processing and actual output skew still require calibration [16].

The resulting architecture is a set of standards-based building blocks plus explicit integration and validation work. IEEE 802.1ASed-2026 supplies a current basis for fault-tolerant timing and time integrity [22], while IEEE P802.1Qee is relevant ongoing work for wireless-aware traffic engineering [23]. A deployment claim should state its measured multicast outage, packet-loss behavior, maximum time error, latency budget, and presentation skew, including the transient during roaming.

14. Glossary

Term	Definition
AC	Access Controller
AP	Access Point
AVB	Audio Video Bridging
AVTP	Audio Video Transport Protocol
BSS	Basic Service Set
CAPWAP	Control and Provisioning of Wireless Access Points
CRF	Clock Reference Format
DMS	Directed Multicast Service
EDCA	Enhanced Distributed Channel Access
FT	Fast BSS Transition
FTM	Fine Timing Measurement
GCR	Groupcast with Retries
gPTP	generalized Precision Time Protocol, the IEEE 802.1AS profile
GTK	Group Temporal Key
IGMP	Internet Group Management Protocol
IGTK	Integrity Group Temporal Key
LMQI	Last Member Query Interval
MLD	Multicast Listener Discovery
MLO	Multi-Link Operation
MSRP	Multiple Stream Registration Protocol
NTP	Network Time Protocol
PIM	Protocol Independent Multicast
PTP	Precision Time Protocol
QI	Query Interval
QRI	Query Response Interval
QRV	Querier Robustness Variable
RRM	Radio Resource Management
STA	IEEE 802.11 station
TAP	Open Compute Project Time Appliances Project
TM	Timing Measurement
TSF	Timing Synchronization Function
TSN	Time-Sensitive Networking
TWT	Target Wake Time
WLAN	Wireless Local Area Network
WMM	Wi-Fi Multimedia
WTP	Wireless Termination Point

Date: October 1, 2026


 This work is licensed under a [Creative Commons Attribution-ShareAlike 4.0 International License](https://creativecommons.org/licenses/by-sa/4.0/).

15. References

- [1] IEEE Std 802.11-2024, IEEE Standard for Information Technology: Telecommunications and Information Exchange between Systems Local and Metropolitan Area Networks: Specific Requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, 2024.
- [2] IEEE Std 802.11ax-2021, Amendment 1: Enhancements for High-Efficiency WLAN, incorporated into IEEE Std 802.11-2024.
- [3] IEEE Std 802.11aa-2012, Amendment 2: MAC Enhancements for Robust Audio Video Streaming, incorporated into IEEE Std 802.11-2024.
- [4] IEEE Std 802.11v-2011, Amendment 8: IEEE 802.11 Wireless Network Management, incorporated into IEEE Std 802.11-2024.
- [5] IEEE Std 802.11r-2008, Amendment 2: Fast Basic Service Set Transition, incorporated into IEEE Std 802.11-2024.
- [6] IEEE Std 802.11k-2008, Amendment 1: Radio Resource Measurement of Wireless LANs, incorporated into IEEE Std 802.11-2024.
- [7] IETF RFC 4541, Considerations for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping Switches, M. Christensen, K. Kimball, F. Solensky, 2006.
- [8] B. Haberman, Ed., "Internet Group Management Protocol, Version 3," STD 100, RFC 9776, March 2025, doi: 10.17487/RFC9776. Obsoletes RFC 3376 and updates RFC 2236.
- [9] IETF RFC 2236, Internet Group Management Protocol, Version 2, W. Fenner, 1997.
- [10] IETF RFC 4605, Internet Group Management Protocol (IGMP)/Multicast Listener Discovery (MLD)-Based Multicast Forwarding ("IGMP/MLD Proxying"), B. Fenner, H. He, B. Haberman, H. Sandick, 2006.
- [11] IEEE Std 802.1BA-2021, IEEE Standard for Local and Metropolitan Area Networks: Audio Video Bridging (AVB) Systems, 2021.
- [12] IEEE Std 802.1AS-2025, IEEE Standard for Local and Metropolitan Area Networks: Timing and Synchronization for Time-Sensitive Applications, 2025.
- [13] IEEE Std 802.1Q-2022, IEEE Standard for Local and Metropolitan Area Networks: Bridges and Bridged Networks, including Forwarding and Queuing Enhancements for Time-Sensitive Streams and the Multiple Stream Registration Protocol.
- [14] Wi-Fi Alliance Technical Committee, WMM Admission Control Technical Task Group, Wi-Fi Multimedia Technical Specification (with WMM-Power Save and WMM-Admission Control), Version 1.2.0, May 7, 2012.
- [15] IEEE Std 802.11-2024, IEEE Standard for Information Technology: Telecommunications and Information Exchange between Systems, Local and Metropolitan Area Networks, Specific Requirements, Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, Clauses 9.6.7.32-9.6.7.33 and 11.21.6 (Fine Timing Measurement frame formats and procedures), 2024.
- [16] IEEE Std 1722-2025, IEEE Standard for a Transport Protocol for Time-Sensitive Applications in Bridged Local Area Networks, 2025.
- [17] IEEE Std 802.11az-2022, Amendment: Enhancements for Next Generation Positioning, including protected ranging, incorporated into IEEE Std 802.11-2024.

Date: October 1, 2026



This work is licensed under a [Creative Commons Attribution-ShareAlike 4.0 International License](https://creativecommons.org/licenses/by-sa/4.0/).

- [18] P. Calhoun, M. Montemurro, and D. Stanley, Eds., “Control And Provisioning of Wireless Access Points (CAPWAP) Protocol Specification,” RFC 5415, March 2009, doi: 10.17487/RFC5415.
- [19] P. Calhoun, M. Montemurro, and D. Stanley, Eds., “Control and Provisioning of Wireless Access Points (CAPWAP) Protocol Binding for IEEE 802.11,” RFC 5416, March 2009, doi: 10.17487/RFC5416.
- [20] H. Zhao, X. Liu, Y. Liu, A. Peter, and M. Sivakumar, “A YANG Data Model for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping,” RFC 9166, January 2022, doi: 10.17487/RFC9166.
- [21] IEEE Std 1588-2019, IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems, 2019.
- [22] IEEE Std 802.1ASed-2026, IEEE Standard for Local and Metropolitan Area Networks: Timing and Synchronization for Time-Sensitive Applications: Amendment 1: Fault-Tolerant Timing with Time Integrity, 2026.
- [23] IEEE P802.1Qee, Draft Standard for Local and Metropolitan Area Networks: Bridges and Bridged Networks: Amendment: Traffic Engineering for Bridged Networks that Include Wireless Technologies, active project, PAR approved September 10, 2025.
- [24] D. L. Mills, J. Martin, J. Burbank, and W. Kasch, “Network Time Protocol Version 4: Protocol and Algorithms Specification,” RFC 5905, June 2010, doi: 10.17487/RFC5905.
- [25] IEEE Std 802.11be-2024, IEEE Standard for Information Technology: Telecommunications and Information Exchange between Systems: Local and Metropolitan Area Networks: Specific Requirements: Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Amendment 2: Enhancements for Extremely High Throughput, 2024.
- [26] Open Compute Project Foundation, OCP Contribution Taxonomy, Draft, Informational, effective June 10, 2026.

Date: October 1, 2026



This work is licensed under a [Creative Commons Attribution-ShareAlike 4.0 International License](https://creativecommons.org/licenses/by-sa/4.0/).

16. License

16.1. Creative Commons Attribution-ShareAlike 4.0 International

OCP encourages participants to share their proposals, specifications and designs with the community. This is to promote openness and encourage continuous and open feedback. It is important to remember that by providing feedback for any such documents, whether in written or verbal form, that the contributor or the contributor's organization grants OCP and its members irrevocable right to use this feedback for any purpose without any further obligation.

It is acknowledged that any such documentation and any ancillary materials that are provided to OCP in connection with this document, including without limitation any white papers, articles, photographs, studies, diagrams, contact information (together, "Materials") are made available under the Creative Commons Attribution-ShareAlike 4.0 International License found here: <https://creativecommons.org/licenses/by-sa/4.0/>, or any later version, and without limiting the foregoing, OCP may make the Materials available under such terms.

As a contributor to this document, all members represent that they have the authority to grant the rights and licenses herein. They further represent and warrant that the Materials do not and will not violate the copyrights or misappropriate the trade secret rights of any third party, including without limitation rights in intellectual property. The contributor(s) also represent that, to the extent the Materials include materials protected by copyright or trade secret rights that are owned or created by any third-party, they have obtained permission for its use consistent with the foregoing. They will provide OCP evidence of such permission upon OCP's request. This document and any "Materials" are published on the respective project's wiki page and are open to the public in accordance with OCP's Bylaws and IP Policy. This can be found at <http://www.opencompute.org/participate/legal-documents/>. If you have any questions, please contact OCP.

Date: October 1, 2026



This work is licensed under a [Creative Commons Attribution-ShareAlike 4.0 International License](https://creativecommons.org/licenses/by-sa/4.0/).

17. About Open Compute Foundation

The Open Compute Project (OCP) brings at-scale innovations and hyperscaler best practices to all, spanning technology domains from the data center to the edge, and the technology stack from silicon, to systems, to site facilities and services. The international OCP Community is made up of organizations and people from hyperscale and tier-2 cloud data center operators, communications providers, colocation providers, diverse enterprises, and technology vendors. The OCP Foundation fosters collaboration within the Community to tackle market challenges in an array of OCP Projects that create open contributions such as specifications, designs, and more. The OCP Solution Provider Program then recognizes the application of those contributions in compliant products, solutions, and data center facilities and services with designations like OCP Inspired, OCP Accepted and OCP Ready. With the tenets of openness, impact, efficiency, scale and sustainability, the OCP engages and educates thousands of engineers every year through many events and webinars. Across many initiatives the OCP Foundation and Community are meeting the market today and shaping the future.

Learn more at: www.opencompute.org.