Array 30 (2026) 100866

Contents lists available at ScienceDirect

Array

journal homepage: www.sciencedirect.com/journal/array

TCA^2: A minimum-assurance time–control framework for GNSS-challenged indoor regions under degraded hybrid links: Integrity-driven arbitration, graceful degradation, and hybrid LF/packet dissemination

Dengfeng Hui [a,b,*], Shamala K. Subramaniam [b], Lili Nurliyana binti Abdullah [b], Abdullah bin Muhammed [b], Hongyan Liu [a], Miao Wang [c], Zijun Wang [c], Yuchun Hui [c], Junyu Zhu [c], Zhongliang Liu [d]

[a] Beijing Wholeway Centron Technology Co., Ltd., Beijing, 100083, China
[b] Universiti Putra Malaysia, 43400, Malaysia
[c] Beijing Normal University Asia-Pacific Experimental School, Beijing, 102209, China
[d] Guangdong Light Industry Technician College, Guangzhou, 511330, China

ARTICLE INFO

Keywords:
Time service
Control assurance
Audit evidence
Service assurance level (SAL)
Minimal control assurance set (MCAS)
LF timecode

ABSTRACT

GNSS-challenged indoor and infrastructure-poor deployments often require not only defensible timekeeping but also a minimum auditable degree of low-rate coordination under degraded links. We present TCA^2, a contract-first framework that makes both time acceptance and minimal control acceptance decidable from operational evidence. Time is accepted when scheduled endpoint observations remain within tolerance θ with availability A(θ) meeting the declared Service Assurance Level (SAL) target; minimal control is accepted when deadline-bounded commands attain R(Δ) within a Minimal Control Assurance Set (MCAS). TCA^2 couples (i) integrity-driven arbitration across GNSS/NTP/LF with explicit holdover and logged transitions, (ii) a timecode-agnostic adapter for national LF profiles, and (iii) hybrid dissemination over one-way LF paths plus an optional bidirectional packet path. In this architecture, LF paths preserve resilient time continuity for legacy or weakest-observability endpoints, whereas deadline-audited control attainment is evaluated only for packet-connected Path C endpoints with ACK/CONFIRM evidence. The evaluation is layered rather than monolithic: DS-1 establishes adapter portability; DS-2 evaluates the time-service contract on 30 legacy display-only clocks in a four-week single-building trial; DS-3 grounds the evidence workflow at field scale through an anonymized footprint of ≈7.85k terminals; and DS-4 demonstrates packet-plane minimal-control evidence in a controlled testbed. In DS-2, audited synchronization improves from ~0.70 under tower-only LF to 0.93–1.00 with local LF augmentation, reaching 1.00 after stabilization; in DS-4, at $\Delta(\text{ACK}) = \Delta(\text{M}) = 30$ s, $R(\text{ACK}) = 0.997$ and $R(\text{M}) = 0.965$ over 593 issued MCAS commands. The contribution lies not in any single underlying carrier or protocol primitive, but in a unified assurance contract that combines multi-source time arbitration, minimum-control preservation, auditable degradation, and offline-recomputable evidence.

1. Introduction

Indoor systems in GNSS-challenged buildings often require not merely nominal synchronization but a defensible, evidence-bound minimum of timing continuity and low-rate control continuity. GNSS, while effective outdoors, may become unreliable indoors because of attenuation, multipath, benign faults, and spoofing [1,2]. In accountable deployments, best-effort behavior is insufficient: operators must be able to determine when a service remains acceptable, what evidence supports that determination, and how the service degrades when evidence is missing, sources disagree, or capacity becomes constrained. We therefore model time and minimal control as auditable services with explicit acceptance contracts. Time is accepted when scheduled endpoint observations remain within tolerance θ with availability A(θ) meeting a declared target, whereas minimal control is accepted when deadline-bounded commands attain R(Δ) on endpoints that can return

* Corresponding author. Beijing Wholeway Centron Technology Co., Ltd., Beijing, 100083, China.
E-mail address: jh@bwt-co.com (D. Hui).

https://doi.org/10.1016/j.array.2026.100866
Received 29 January 2026; Received in revised form 18 April 2026; Accepted 28 April 2026
Available online 8 May 2026

completion evidence. The contract makes acceptance decidable by fixing the measurand, the uncertainty treatment, and the missing-evidence rule in a manner consistent with metrological decision practice [3,4]. Acceptance is expressed through SAL-2/1/0, while control obligations are deliberately narrowed to a Minimal Control Assurance Set (MCAS) that must remain preservable under degradation. MCAS denotes the smallest semantically essential set of control instructions and status messages that the system continues to protect when integrity or capacity deteriorates, thereby preserving a minimum auditable service core rather than an undefined best-effort behavior.

Implementing such contracts in GNSS-challenged buildings introduces four coupled tensions: integrity versus availability; path diversity versus semantic continuity; coverage versus portability, because LF penetrates buildings but national formats vary [5–9]; and control versus capacity, because packet-plane telemetry and command evidence remain airtime-limited in constrained deployments [10–15]. TCA^2 addresses these tensions through integrity-driven arbitration across GNSS/NTP/LF, a timecode-agnostic adapter for national LF profiles, and a hybrid dissemination architecture in which LF paths preserve resilient time continuity and coarse trigger dissemination, while the packet plane carries ACK/CONFIRM-backed minimal-control evidence. When NTP is used, its contribution can be protected by NTS [16,17]. Under packet-capacity or connectivity loss, the system narrows its contractual claim through SAL/MCAS rather than asserting unchanged control guarantees on one-way paths.

Contributions. This paper makes four contributions. (1) A minimum-assurance service contract: endpoint-defined acceptance metrics for time and minimal control, A(θ) and R(Δ), tied to SAL and MCAS, together with conservative missing-evidence rules and a minimum log schema for offline recomputation. (2) An integrity-driven hybrid architecture: explainable multi-source arbitration with explicit holdover, a timecode-agnostic adapter for national LF profiles, and path-separated dissemination over one-way LF and an optional bidirectional packet plane. (3) Auditable graceful degradation semantics: SAL/MCAS policies that preserve TIME_STATE and MCAS traffic under integrity loss or airtime pressure while deferring non-MCAS traffic. (4) A layered evidence design: DS-1 establishes adapter portability; DS-2 evaluates the time-service contract on weakest-observability legacy endpoints; DS-3 grounds the evidence workflow at field scale; and DS-4 demonstrates packet-plane minimal-control attainment. We do not claim deadline-audited control attainment for one-way legacy LF paths.

2. Related work

Time synchronization spans diverse domains. In the Internet, NTP provides ubiquitous best-effort synchronization but its accuracy depends on path symmetry and delay variation, so any “synchronized” claim must specify the endpoint and measurement method [16]. For tighter bounds in controlled networks, IEEE 1588 Precision Time Protocol (PTP) and its Ethernet profile (gPTP under IEEE 802.1AS) achieve sub-millisecond precision using hardware timestamps and deterministic networking assumptions [18,19]. In industrial IoT settings, surveys discuss practical PTP deployments, TSN integration, and remaining challenges for robust clock synchronization [20].

Timing robustness and security have been widely studied. For example, NTP's basic design is vulnerable to on-path attacks, leading to the development of Network Time Security (NTS) which adds cryptographic authentication [17]. Similarly, concerns over GNSS spoofing and other deceptive timing attacks [1,2] have driven research into detection and mitigation, highlighting that a dependable system must prioritize integrity and explainability of time sources - not just precision. Collectively, prior work indicates that a timing service should provide auditable evidence of when it is synchronized, when/how long it holds over during source outages, and why it switches references.

National LF time broadcasts (e.g., WWVB, DCF77, JJY, BPC) are one of the few technologies that penetrate buildings at scale with low-cost, low-power receivers [5–9]. However, indoor reception quality is highly site-dependent and the timecode formats vary by country, so historically such LF solutions have been treated as fixed to their region rather than leveraged via a portable abstraction.

Moreover, metrology standards emphasize explicit decision rules that account for measurement uncertainty, providing a principled template for making acceptance decisions auditable rather than relying on ad hoc thresholds [21].

Low-Power WANs like LoRaWAN provide low-energy connectivity for periodic telemetry and control, but they impose constraints on downlink capacity due to airtime limits and regulatory duty cycles [10]. LoRaWAN does include basic time-sync features (e.g., DeviceTime requests and Class B beacons) and an application-layer clock sync protocol aiming for ±1 s accuracy [12], but its scalability is limited by collisions and downlink bottlenecks in dense deployments [13–15]. These findings motivate our design choice to enforce a minimal guaranteed control set (MCAS) and to implement scheduling with explicit admission control and prioritization, ensuring deadline-bound control messages are delivered even under congestion.

Operational assurance often relies on telemetry and anomaly detection at scale [22,23]. For legacy, display-only endpoints, vision-based reading of analog clocks can serve as an independent observation channel with quantifiable error [24].

Recent open-access work in Array explores learning-based multipath detection for GNSS signals, complementing the spoofing-focused literature and reinforcing the need for explicit evidence when GNSS is used as an input source [25].

Gap and positioning: Prior studies have typically treated time transfer accuracy, LPWAN capacity, source integrity, and operational monitoring as separate concerns. TCA^2 does not claim to invent each constituent mechanism in isolation. Rather, its novelty lies in unifying multi-source time arbitration, minimum-control preservation, auditable degradation, and offline-recomputable evidence into a single assurance contract that remains interpretable across heterogeneous endpoints and degraded hybrid links. This distinction is important: the present paper is not a conventional indoor clock-synchronization study. Legacy clocks are used only as a weakest-observability endpoint class for stressing the acceptance interface, whereas deadline-audited minimal-control attainment is demonstrated on packet-connected endpoints.

3. TCA^2 service model and acceptance metrics

3.1. Two coupled services: time and control

TCA^2 models an indoor timing system as two coupled services-time and control-each with acceptance criteria that are decidable from operational logs. Both services are defined at the acceptance interface (what an endpoint displays, timestamps, or executes) rather than at internal protocol states, so that compliance can be audited and, when needed, independently verified.

Time service. The service core maintains a reference time on a common time scale by fusing multiple sources (e.g., GNSS, NTP, national LF broadcasts). Let $\boldsymbol{T_E(t)}$ be the time indicated by endpoint $\boldsymbol{E}$ (e.g., a clock or controller) at true time $\boldsymbol{t}$; define the endpoint's time error as $\boldsymbol{\Delta_E(t) = T_E(t) - t}$. A time service contract specifies a tolerance $\boldsymbol{\theta}$ and accepts the endpoint if $|\boldsymbol{\Delta_E(t)}| \leq \boldsymbol{\theta}$ for a sufficient fraction of sample points over a given audit window (Section 3.2). This contract is defined relative to the system's internal reference $\boldsymbol{T}$; traceability to external standards (e.g., UTC) depends on which sources are used and is handled separately as an external consistency issue.

Control service. The service plane issues commands (and optional schedules) to packet-connected endpoints on the packet plane (Path C) and, where applicable, to site-local controllers such as a rebroadcast/adapter module. For each auditable command instance, the server records the issue time and the server-observed completion-evidence time (ACK/CONFIRM) on a common timebase. One-way legacy LF endpoints

on Paths A/B provide no native uplink acknowledgements and typically no device logs; accordingly, deadline-bounded command attainment R (Δ) is evaluated only where packet-plane completion evidence exists and is not claimed for DS-2 or DS-3. When packet connectivity degrades, TCA[2] may continue to disseminate time continuity or coarse trigger semantics over LF, but it does not preserve the same command-attainment claim unless the required completion evidence remains available. Manual or vision-based observations, when used, are auxiliary evidence channels for maintenance or post hoc corroboration and do not replace packet-plane completion evidence for R(Δ).

3.2. Acceptance metrics

(M1) Endpoint time error (measurand). For a given endpoint E, define its time error as Δ_E(t) = T_E(t) − T_ref(t), where T_E(t) is the endpoint-indicated time at real time t and T_ref(t) is the system reference time produced by the arbitration core. In metrological terms, Δ_E is the measurand, and each retained observation should be accompanied by an explicit uncertainty statement [3,4]. Depending on endpoint class, Δ_E may be obtained either from direct digital telemetry or from an observation channel for legacy devices, including manual or vision-assisted reading of a display-only interface. What matters contractually is not the acquisition modality itself, but whether the retained evidence specifies the observation method, sampling instant (or window), derived offset, and associated uncertainty bound in a form suitable for offline recomputation.

In either case, each error sample is paired with a conservative observation-uncertainty bound u_e, archived as evidence for offline recomputation. We budget u_e from (i) display resolution/quantization u_res, (ii) readout uncertainty u_read (human or vision interpretation), and (iii) audit-time alignment uncertainty u_align between the scheduled sample instant and the effective read/capture time; a simple reproducible form is:

$$u_e = \sqrt{u_{res}^2 + u_{read}^2 + u_{align}^2} \tag{1}$$

Note that mechanism drift or actuation imperfections manifest in the measured endpoint error ΔE itself; u_e only bounds how accurately that realized behavior is observed.

(M2) Time availability within threshold. We measure time availability via scheduled sampling. Each endpoint E is sampled periodically (interval τ_s, e.g., 60 s), yielding either a time error $\Delta_E(t_k)$ (with uncertainty) or a missing observation (if no data are received or the reading is unreadable at that sample time). A sample at time t_k is considered available if $|\Delta E(t_k)| + u_e(t_k) \le \theta$, where $u_e(t_k)$ is the observation-uncertainty bound associated with that sample (Section 3.2 (M1)). Any scheduled sample that is missing (or invalid) is counted as unavailable unless an independent trusted channel corroborates the endpoint's time for that instant. A(θ) is evaluated on a scheduled evidence set $\mathcal{T} = \{t_1, \ldots, t_n\}$ defined by the service contract and constrained by the endpoint's observable interface. For instrumented endpoints $\mathcal{T}$ may be dense (trace-like), while for legacy endpoints $\mathcal{T}$ corresponds to periodic audit instants; missing scheduled evidence is counted as unavailable. Over an audit window containing N scheduled samples for endpoint E, we compute the endpoint's availability as

$$A_E(\theta) = \frac{1}{N}\sum_{i=1}^{N} \mathbf{1}\{|\Delta_E(t_i)| + u_e(t_i) \le \theta\} \tag{2a}$$

where the indicator $\mathbf{1}\{\cdot\}$ equals 1 for each available sample and 0 otherwise. For site-level acceptance, we aggregate by taking a low percentile (5th percentile) of the endpoints' individual availabilities $\{A_E(\theta)\}$, so that

$$A_{site}(\theta) = P_5(\{A_E(\theta)\}_{E \in site}) \tag{2b}$$

(Lower percentiles emphasize worst-case endpoints, while higher moments can be reported separately for diagnostic purposes.) Fig. 1 illustrates scheduled sampling and how $A_E(\theta)$ is derived for one endpoint and then aggregated to $A_{site}(\theta)$. *(The figure is schematic and does not use experimental data; it is included only to illustrate the computation pipeline.)*

(M3) Control attainment within deadline. For a given command class with a deadline requirement Δ, consider all command instances issued within the audit window. Each instance has an issue time t_i (when the server transmits the command) and a server-observed completion-

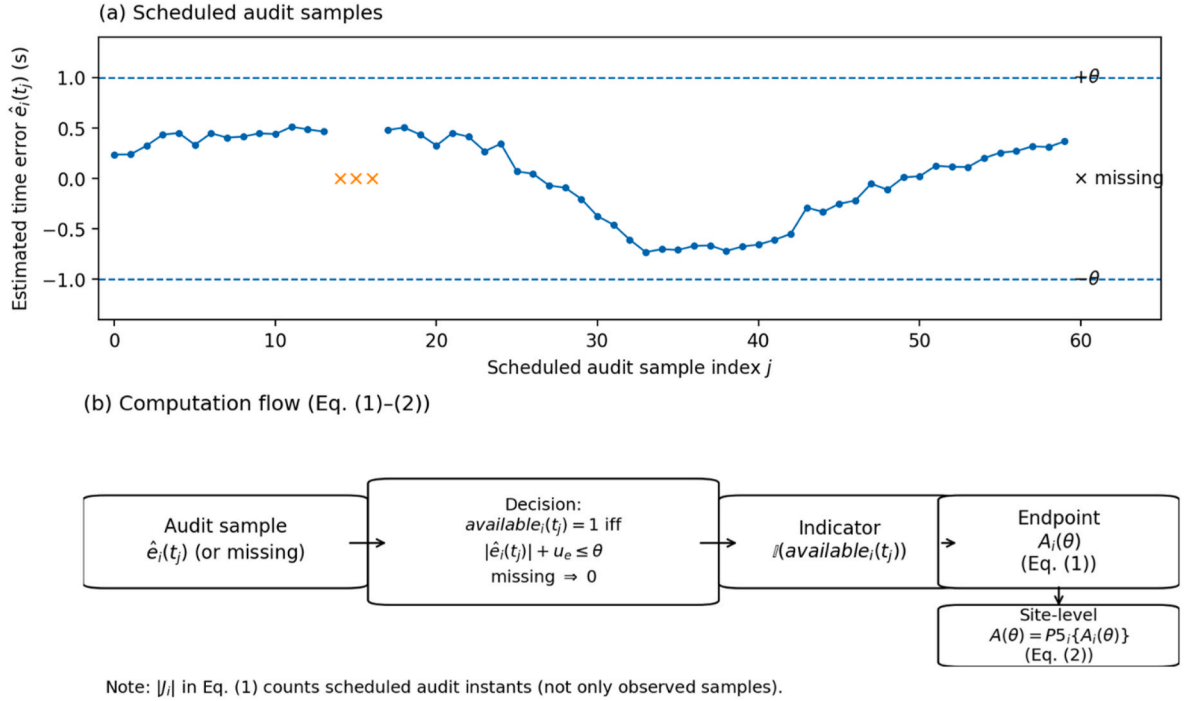

Fig. 1. Sampling-based time-service acceptance illustration. *(a) Scheduled endpoint time-error samples with tolerance band ±θ (shaded green region shows acceptable error range). (b) Flowchart for computing endpoint availability A(θ) and site-level A_site(θ) from audit samples.* (Illustrative example; not actual field data.)

evidence time t_c (e.g., when an ACK/CONFIRM is received), both on a common time base. The auditable delivery latency for instance j is $\Delta_{\mathrm{ctrl},j} = t_{c,j} - t_{i,j}$. We define the command attainment ratio as the fraction of audited instances that meet the deadline:

$$R(\Delta) = \frac{\#\{j \in \Omega : \Delta_{\mathrm{ctrl},j} \leq \Delta\}}{|\Omega|} \tag{3}$$

Equivalently,

$$R(\Delta) = \frac{1}{|\Omega|} \sum_{j \in \Omega} \mathbf{1}\{t_{c,j} - t_{i,j} \leq \Delta\} \tag{4}$$

where Ω is the set of issued command instances in the window, excluding those legitimately cancelled or aborted as evidenced by logs. If completion evidence is missing for an instance, it is counted as not attained (i.e., the indicator $\mathbf{1}\{\cdot\} = 0$). For broadcast commands sent to multiple endpoints, Ω is expanded over (command, endpoint) pairs, so that each addressed endpoint must individually satisfy the deadline. We compute $R(\Delta)$ separately for MCAS and non-MCAS traffic to keep degradation behavior auditable and to ensure that load-induced deferrals of non-MCAS traffic do not obscure the guaranteed attainment of MCAS commands.

3.3. Service Assurance Levels (SAL) and auditability

Service Assurance Levels (SAL). To make acceptance operational and auditable, TCA^2 defines three contractual states-SAL-2/1/0-computed from the acceptance metrics in Section 3.2 over an audit window W. Each SAL specifies (i) threshold requirements for time availability and control attainment, (ii) the minimum evidence needed to recompute them, and (iii) explicit, logged triggers for degradation and recovery (Table 1, Fig. 2). Missing required evidence is treated conservatively as the corresponding criterion not satisfied unless independently corroborated.

We distinguish the run-time SAL status, SAL_{state}, from the acceptance outcome, SAL_{accept}, reported over an audit window (e.g., per session or per week). The run-time status may be adjusted on short decision windows (e.g., minute-scale) based on ongoing integrity screening and the airtime-based capacity indicator U_W. By contrast, SAL_{accept} is evaluated post hoc from the complete archived evidence set for that window. Thus, SAL_{state} can downgrade promptly when integrity is lost or capacity is stressed, whereas the conclusive SAL_{accept} is determined after the fact from the full evidence record of the period.

Table 1 summarizes the acceptance clauses and minimum audit evidence for each SAL, and Fig. 2 illustrates the corresponding run-time transitions.

Table 1
Service Assurance Levels (SAL) and acceptance clauses (A(θ) denotes site-level P5 endpoint availability; Section 3.2).

SAL	Time service acceptance	Control service acceptance	Auditability/evidence
SAL-2	A(θ) ≥ 0.99	R_MCAS(Δ_M) ≥ 0.99; R_non(Δ_N) ≥ 0.95	Integrity arbitration enabled; switching/holdover reasons logged; missing samples counted as unavailable.
SAL-1	A(θ) ≥ 0.95	R_MCAS(Δ_M) ≥ 0.95; non-MCAS may be deferred	Integrity arbitration enabled; MCAS prioritized; explicit SAL transition events logged.
SAL-0	Best effort	Best effort	Minimal operational logging; no contractual acceptance.

Note: Contract parameters (defaults unless noted): θ = 1 s; Δ_M = 30 s; Δ_N = 120 s; Δ_ACK = 30 s; τ_int = θ; W_switch = 120 s; W_recover = 600 s; τ_s = 60 s (instrumented endpoints); capacity high/low watermarks U_high, U_low (site-configurable; used for SAL transitions and logged).

3.4. Integrity as first-class: arbitration, holdover, and explainable switching

TCA^2 treats time integrity as a first-class requirement: anomalies must be detectable, source switching must be explainable, and continuity must be maintained via an explicit holdover mode when no validated reference is available. The time arbitration core outputs the reference time T_{ref} together with an integrity state (HEALTHY, SUSPECT, or HOLDOVER) and logs every switching/holdover event with a reason code and supporting evidence (see Table 1).

At each decision point, the core collects candidate observations from all available sources (e.g., GNSS, NTP, and national LF timecodes), normalizes them onto a common seconds time base, and performs consistency checks. Each source provides a time estimate and auditable quality indicators (e.g., GNSS lock status, NTP stratum/reachability and leap flags, or LF decode confidence), which are recorded as part of the switching evidence. Sources failing basic validity screening are excluded. For the remaining sources, the core computes pairwise offsets and flags a source as SUSPECT when its deviation exceeds the integrity tolerance τ_{int} for at least the switching window W_{switch}. The reference time T_{ref} is then derived from the remaining consistent candidate(s) using a deterministic, replayable selection policy, and the chosen candidate(s) and rationale are logged for audit.

If no candidate remains mutually consistent (e.g., inputs are lost or disagree), the system enters HOLDOVER: T_{ref} advances monotonically using a local oscillator/RTC, bounded-accuracy claims are suspended, and the holdover entry time, the last-good synchronization time, and a conservative drift bound are recorded. Recovery is staged to prevent oscillation: a candidate must remain consistent for W_{recover} before exiting HOLDOVER, and any step or slew applied during re-synchronization is logged as an integrity event.

Throughout, audit timestamps and command latencies are maintained on a monotonic clock, while UTC/civil-time rendering is applied only at presentation. Any mapping discontinuities (e.g., display jumps on re-lock) are treated as explicit, logged measurement context rather than hidden state. Fig. 3 summarizes the arbitration workflow and integrity-state transitions for explainable switching and holdover.

3.5. Minimal Control Assurance Set (MCAS)

MCAS is the minimal set of control commands and status messages that the system guarantees to uphold even under degraded conditions (e.g., in post-disaster or remote-edge scenarios when switching from SAL-2 down to SAL-1). This set is defined in terms of system semantics and kept deliberately small. Each MCAS message type has a specific role, a fixed or bounded payload size, and a deadline or freshness requirement that is subject to auditing. Table 2 lists the MCAS message types and the minimum log fields required for offline audit of each. In essence, MCAS defines the critical commands and beacons that must always get through or be accounted for, whereas non-MCAS traffic can be postponed if needed. For example, in a degraded remote-operation scenario, MCAS may be limited to safe-state transitions, emergency-stop semantics, or minimal schedule/trigger dissemination, ensuring that a device can still receive and evidence the smallest contractually protected instruction set even when richer telemetry is delayed or suppressed. Likewise, in indoor infrastructure management, MCAS may consist of minimal scheduling triggers, aborts, and confirmations required for safe operation, while non-essential updates are deferred. These examples illustrate low-rate critical continuity rather than high-frequency continuous closed-loop control.

When network capacity is under stress (airtime congestion), the scheduler triggers a controlled degradation from SAL-2 to SAL-1 before the system is overwhelmed, in order to reserve airtime for MCAS messages and delay non-essential traffic (We detail this mechanism in Section 5.).

The **TIME_STATE** beacon periodically reports a coarse reference time

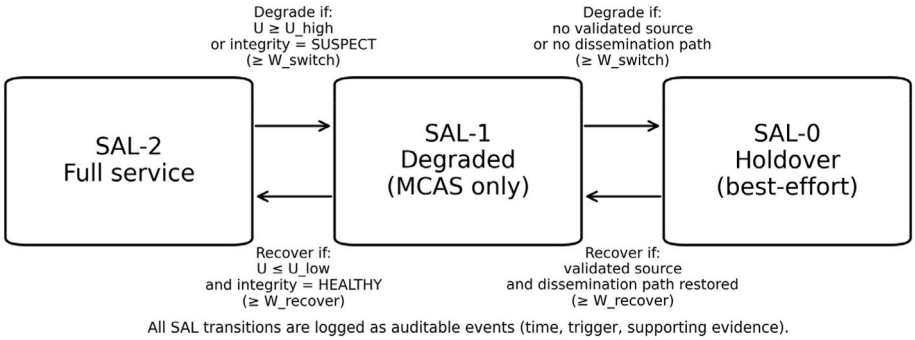

Fig. 2. Service Assurance Level (SAL) run-time transitions with auditable triggers and hysteresis windows $\boldsymbol{W_{switch}}$ and $\boldsymbol{W_{recover}}$ (illustrative).

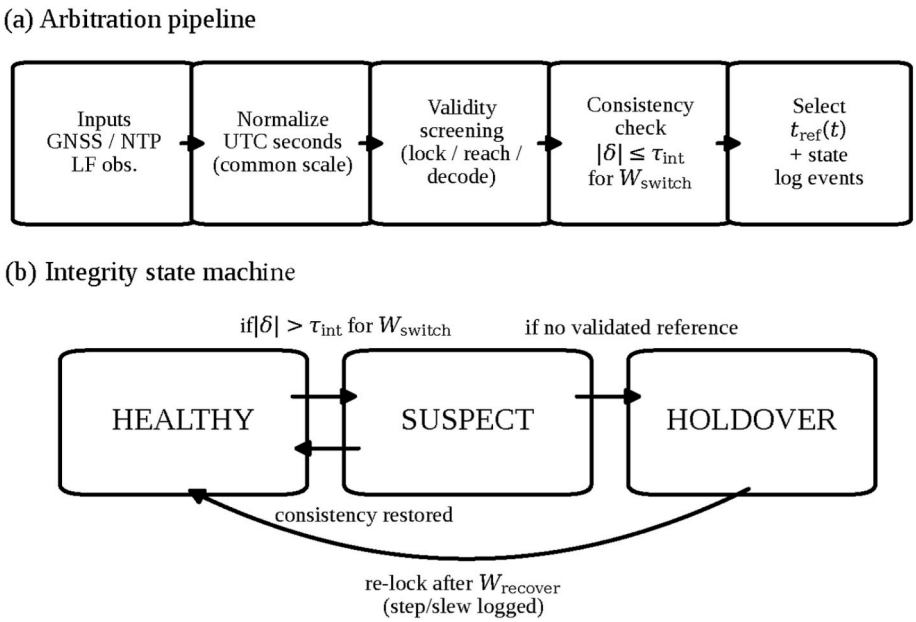

Fig. 3. Integrity-driven source arbitration with explainable switching and holdover. Candidate observations (GNSS/NTP/LF) are normalized to a common seconds scale, screened by validity checks, and compared by cross-source consistency $\delta(s, s')$ against τ_{int} over W_{switch}. The arbiter outputs $t_{ref}(t)$ together with an integrity state (HEALTHY/SUSPECT/HOLDOVER); all switch/holdover events are logged with a reason code and supporting evidence.

plus the current integrity/holdover status; its acceptance requirement is a freshness bound (τ_s) rather than a delivery deadline. **EXECUTE_AT** schedules an action at a specified reference time T_{exec}, and **ABORT** cancels a pending action - both are MCAS commands that must be delivered by their deadlines Δ_M. **CONFIRM** is sent from endpoint to server as completion evidence (acknowledgment or execution report); it must arrive within Δ_{ACK} and can include an actual execution timestamp for audit. The contract ties the TIME_STATE interval and control deadlines together (e.g., ensuring an ACK/CONFIRM can be piggybacked on the next beacon without exceeding Δ_{ACK}), so that acknowledgment latency remains bounded and auditable.

If a required MCAS confirmation is missing, we treat it as a failure (violation) in the audit, per the policy of Section 3.2. Meanwhile, non-MCAS traffic (e.g., bulk telemetry or high-rate samples) is only sent when there is slack capacity, and in SAL-1 it may be buffered or downsampled to reserve airtime for MCAS (Section 5 details how the scheduler enforces this partition under LPWAN constraints.).

4. System architecture

This section describes the TCA^2 architecture: a service plane (arbitration/holdover, timecode adaptation, scheduling, and evidence logging) plus three parallel dissemination paths. The contract is transport-agnostic: the same SAL/MCAS logic applies whether time/control are carried by tower LF, site-local LF, GNSS/NTP inputs, or a packet plane (LPWAN or LAN). Fig. 4 summarizes the end-to-end design.

4.1. Overview and dissemination paths

Overview of dissemination paths. TCA^2 supports three dissemination paths (A, B, C) alongside the core service plane that generates the reference time, integrity state, and audit evidence for the acceptance metrics (Section 3). Each path is an independent realization of the same service contract at the endpoint interface - the deployment can use any or all paths in combination rather than a one-size-fits-all stack.

Path A – Tower LF broadcast to legacy clocks. Path A uses direct reception of a national LF timecode by radio-controlled clocks. This approach offers wide-area coverage with minimal local infrastructure. However, indoor reception quality is highly site-dependent and may suffer from blind spots; moreover, the legacy clocks typically do not report any signal quality or error, which limits the evidence available for audit.

Path B – Timecode adapter with local LF rebroadcast. Path B improves indoor coverage by **rebroadcasting** the reference time as a local LF signal that matches the national format expected by legacy clocks. This explicit adapter boundary allows the encoding process to be

Table 2
MCAS message types and required evidence fields (service-level contract).

Message	Direction	Purpose	Deadline/ freshness	Required log fields (minimum)
TIME_STATE	server → endpoint	Periodic time + integrity/ holdover beacon (coarse t_{ref} + state)	freshness $\leq \tau_s$	send_ts, endpoint_id (or group_id), ref_seq, t_ref_sec, integrity_state, holdover_flag
EXECUTE_AT	server → endpoint	Schedule a time-aligned action at reference time T_{exec}	Δ_M (MCAS)	cmd_id, endpoint_id, send_ts, T_exec, action_code, priority = MCAS
ABORT	server → endpoint	Cancel a pending action/ enter a safe-stop state	Δ_M (MCAS)	cmd_id, endpoint_id, send_ts, target_cmd_id, reason_code, priority = MCAS
CONFIRM	endpoint → server	Completion evidence for attainment and auditing (receipt/ optional execution confirm)	Δ_{ACK}	cmd_id, endpoint_id, recv_ts (server timebase), status_code, [t_event]

Note: τ_s, Δ_M, and Δ_{ACK} are contract parameters; missing required evidence is treated as the corresponding criterion not satisfied unless independently corroborated (Section 3.2).

verified offline (via conformance tests and replay), decoupling the timecode format conversion from the physical radio environment. Bench-level prototype evidence and representative LF endpoints for Path B are documented in Appendix A2 (non-normative).

Path C – Packet network (LPWAN or LAN) for digital endpoints. Path C provides a bidirectional **packet plane** to distribute time beacons, send MCAS control commands, and collect evidence/telemetry from digital endpoints. It can be realized over a constrained LPWAN link (e.g., LoRaWAN) or an unconstrained IP network (e.g., Wi-Fi/LAN) depending on deployment conditions. When operating over a constrained link, Path C's capacity is limited by airtime/duty cycle, so the scheduler must enforce SAL-driven priority (ensuring MCAS messages are delivered first under load). Unlike Path C, Paths A and B are one-way time broadcasts and do not support command acknowledgments (R(Δ)), which limits their control capabilities and auditable scope.

To make the bidirectional control interaction on Path C explicit, Fig. 5 sketches the control-command downlink and the status/ acknowledgment uplink under the dual-channel design, and illustrates one-to-many extensibility.

Control scope across paths. Paths A/B target one-way legacy or weakest-observability endpoints and therefore do not, by themselves, provide native completion evidence for deadline-audited command attainment. Accordingly, R(Δ) is asserted only for packet-connected endpoints on Path C. Any vendor-specific LF side channels are treated as optional hints and are not assumed in the contract.

Path-transition semantics under degradation. When a deployment temporarily loses packet-plane capacity or connectivity, TCA^2 does not transfer the full R(Δ) guarantee from Path C onto Paths A/B. Instead, the contractual claim narrows: LF dissemination may preserve time continuity, TIME_STATE freshness, or coarse trigger semantics, while deadline-audited command attainment remains asserted only for those endpoints and intervals with retained packet-plane completion evidence. Any such narrowing is reflected by SAL transition logs and corresponding acceptance outcomes. We evaluate Paths A/B via DS-2/DS-3, adapter portability via DS-1, and packet-plane MCAS evidence via DS-4.

4.2. Timecode-agnostic adapter and portability validation

Timecode-agnostic adapter. National LF timecode signals (WWVB, DCF77, JJY, BPC, etc.) all encode local civil time in their own format (BCD fields, parity bits, DST/leap indicators, modulation patterns, etc.) [5–9]. In TCA^2, a Timecode Adapter module hides these differences by converting the country-specific timecode into a standard internal format. The adapter normalizes each received timestamp to the universal UTC seconds scale for use in the core (arbitration, scheduling, logging), but it preserves the local calendar info (DST flags, leap-second announcements, etc.) for display if needed.

On the output side, the adapter takes the system reference time (plus the current integrity state and any leap-second or DST flags) and encodes it into the specific national timecode format for local broadcast. Because the rest of the system works only with the internal normalized time, swapping from one timecode profile (e.g., BPC) to another (e.g., DCF77 or JJY) is just a profile change - it does not affect the SAL logic or acceptance metrics at all. To verify such portability explicitly (rather than assume it), we provide an offline conformance & replay test harness (Fig. 4(b)).

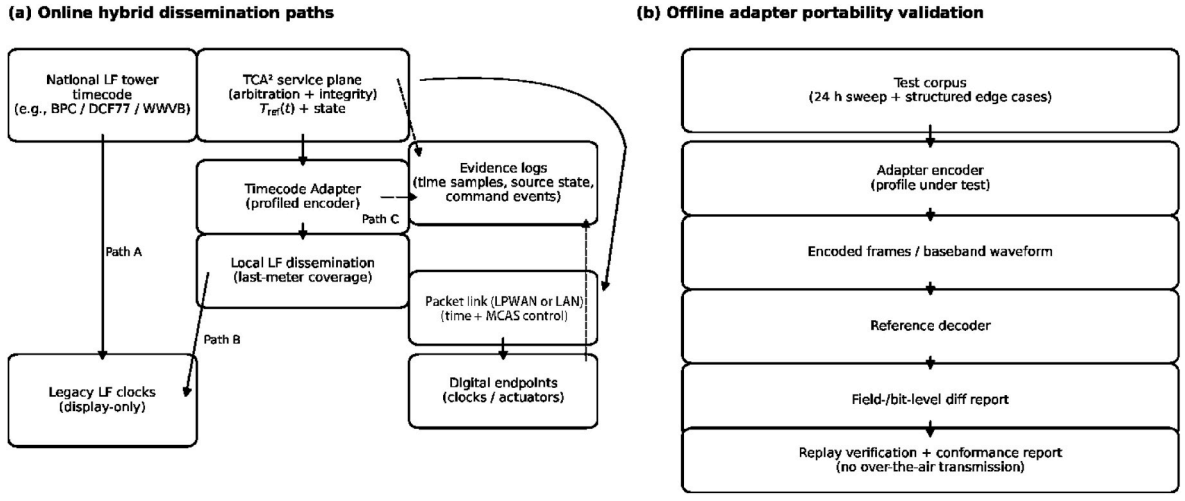

Fig. 4. Hybrid dissemination and portability evidence. (a) Three paths: A tower LF→legacy clocks; B site-local LF rebroadcast via the Timecode Adapter to close blind spots; C packet plane (LPWAN/LAN) for time + MCAS to digital endpoints. Dashed arrows indicate evidence logging. (b) Offline adapter conformance/replay: corpus→encoder→independent decode→field/bit diffs (no RF).

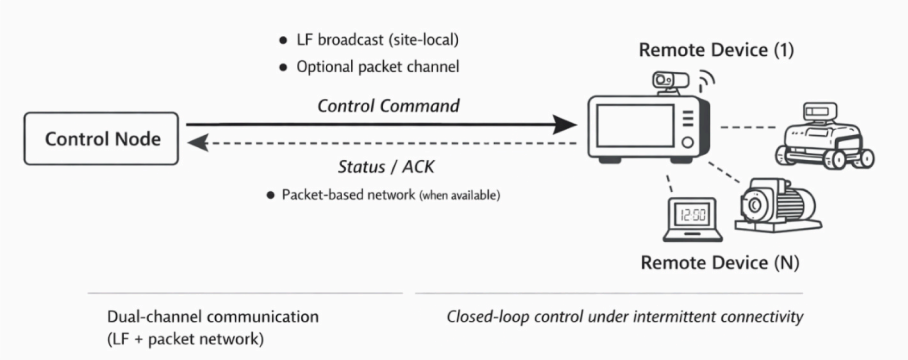

Fig. 5. TCA^2 dual-channel control and feedback loop with one-to-many extensibility (illustrative). *Control commands are disseminated* via *site-local LF broadcast and/or an optional packet plane implementing MCAS. Remote devices return status updates and acknowledgments* via *the packet-based network when available, enabling closed-loop control under intermittent connectivity.*

The validation harness operates in two modes. In conformance mode, it runs the adapter on a known set of test timestamps (e.g., a 24-h minute-by-minute corpus plus edge cases) and then independently decodes the adapter's outputs to check that every field matches expectation (reporting any bit-level differences). In replay mode, it takes recorded or simulated baseband waveforms of timecode signals and feeds them through a standard LF receiver/decoder, verifying that the adapter's output can drive a real receiver correctly end-to-end without on-air transmission. Table 3 summarizes the test corpus used, the pass criteria, and the key log fields recorded for audit.

4.3. Compliance-gated local LF augmentation

Deployments can optionally use a low-power LF rebroadcaster (e.g. a USB-stick LF transmitter) to eliminate indoor blind spots for legacy clocks (Path B's last-meter coverage). We treat this local LF augmentation as a strictly optional, compliance-gated feature rather than a default: enable/disable actions are explicitly controlled and logged, and any on-site use must adhere to safety and radio regulations. If local rebroadcast is not permitted in a region, the module simply remains disabled, and the achievable service level will be limited by direct coverage (to be reported accordingly).

For sites with severe attenuation, an optional building-scale LF augmentation design can be commissioned, but it remains strictly compliance-gated and auditable. In particular, the commissioning view is not a deployment prescription: optional rooftop injection points, preferred coupling paths along building PE or steel frameworks, and dominant loss paths are site-dependent engineering choices recorded as commissioning metadata rather than assumed system constants. Their purpose is to close last-meter coverage in obstructed buildings while keeping the service claim explainable: the enable/disable state, operating profile, and relevant commissioning metadata are logged, and achievable acceptance remains governed by the same A(θ)/SAL criteria. The corresponding schematic and commissioning definitions are provided in Appendix A.

In summary, Sections 4.1–4.3 have outlined a timecode-agnostic multipath dissemination architecture, in which the service interface (contract) stays consistent across different carriers and national timecode profiles. Next, Section 5 describes how we enforce this contract under LPWAN capacity constraints via airtime-aware scheduling and an observation layer - ensuring the system degrades in a controlled manner while still meeting MCAS obligations.

5. Airtime-aware scheduling and observation layer

This section makes the contract enforceable on constrained packet links. We describe (i) an airtime-aware scheduler that prioritizes TIME_STATE and MCAS traffic and triggers SAL transitions before deadlines are missed, and (ii) an observation layer that logs the minimal evidence needed to recompute acceptance and explain each transition offline. The same logic applies on unconstrained LANs; airtime triggers simply become inactive.

5.1. Airtime accounting, duty-cycle constraints, and congestion triggers

In LPWANs, effective capacity is governed by time-on-air (ToA) and channel-occupation constraints, so airtime is often a more actionable quantity than nominal bitrate. TCA^2 therefore maintains an airtime ledger for every scheduled downlink and observed uplink transmission, computed over a sliding observation window T_{obs}. For a message m sent with PHY profile p (e.g., SF/BW/CR) and payload length $L(m)$, the scheduler estimates its time-on-air $\mathrm{ToA}(m,p)$ using a stated ToA model, and records the estimate together with the message class (MCAS vs. non-MCAS) and timestamps. Windowed airtime utilization is computed as

$$u(t;T_{\mathrm{obs}}) = \frac{\sum_{m\in\mathscr{P}(t,T_{\mathrm{obs}})} \mathrm{ToA}(m,p_m)}{T_{\mathrm{obs}}} \tag{5}$$

where $\mathscr{P}(t,T_{\mathrm{obs}})$ enumerates transmissions within the window.

Under duty-cycle regimes, u matches the usual "on-time over observation period" accounting; in non-duty-cycle regimes it remains a consistent proxy for congestion risk because it reflects how much airtime the system attempts to occupy. We treat the high/low watermarks U_{high}

Table 3
Offline validation of timecode-adapter profiles (conformance and replay).

Mode	What it verifies	Default corpus (example)	Pass criterion	Minimum logged fields
Conformance	Field-level semantic correctness via independent decode	24 h sweep (1440 min frames) + structured edge cases (year boundary, leap day, DST enter/exit if applicable, leap-second flag, etc.)	All decoded fields match inputs; no invalid frames in the corpus	adapter_profile_id; adapter_version; input_time_tag; encoded_frame_hash; decoder_version; diff_summary
Replay	Waveform-level timing/ framing through a receiver/ decoder chain (no RF)	Same corpus rendered as recorded/ synthesized baseband waveform	Decoder lock and correct field recovery within profile-defined tolerance for all test vectors	adapter_profile_id; adapter_version; waveform_hash; replay_chain_id; decoder_version; error counters

Note: The corpus and tolerances are profile-specific; logged hashes enable exact replay and post-hoc diagnosis.

and U_{low} ($U_{\text{low}} < U_{\text{high}}$) as explicit configuration parameters chosen to respect regional constraints and to provide hysteresis against oscillation.

For capacity planning, an approximate upper bound on supported endpoints is:

$$N_{\max} \approx U \Big/ \sum\nolimits_j \left(\text{ToA}_j / T_j\right) \tag{6}$$

which provides a transparent capacity envelope for MCAS budgeting.

SAL transitions are policy-driven and auditable. A SAL-2→SAL-1 downgrade is triggered when sustained airtime pressure threatens MCAS delivery, i.e., (i) $U_W > U_{high}$ for at least W_{switch}, (ii) the scheduler predicts that queued MCAS traffic cannot meet Δ_M or Δ_{ACK} under the remaining airtime budget (including bounded retransmissions), or (iii) integrity events reduce the set of validated references or dissemination paths. Recovery (SAL-1→SAL-2) applies hysteresis: $U_W < U_{low}$ for $W_{recover}$ and MCAS attainment meets the SAL-2 target over the same interval while the arbiter remains in a stable integrity state. Each transition is logged with a reason code and minimal supporting counters (utilization summary, queue depth, missed-deadline counts, integrity state) to enable offline reconstruction.

In dense sites, the scheduler treats TIME_STATE fan-out and ACK/CONFIRM fan-in as first-class budget items rather than incidental traffic. To prevent uplink collapse, reserved airtime is assigned to MCAS confirmations before admitting non-MCAS telemetry, and TIME_STATE dissemination may be grouped where contractually permitted. Under sustained pressure, the system first throttles or buffers non-MCAS traffic, then narrows observation density for non-essential flows, and only then downgrades SAL once predicted MCAS deadlines or confirmation freshness can no longer be maintained within the remaining budget. The design intent is predictive degradation: SAL-2→SAL-1 occurs before contractual MCAS failure becomes systemic, not after queue saturation has already invalidated the claim. DS-3 is used only to ground the plausibility of dense operational contexts and not as a controlled congestion benchmark.

5.2. Priority queues and degradation triggers

The scheduler maintains separate logical queues for time beacons, MCAS traffic, and non-MCAS traffic, and enforces a strict priority order: time > MCAS > non-MCAS. This separation makes the degradation behavior policy-driven rather than an incidental consequence of queueing, ensuring that best-effort traffic cannot consume the slack required for deadline-bounded control. Within the MCAS class, packets are served in deadline order (e.g., earliest-deadline-first) with bounded retries, while non-MCAS packets are admitted only when the airtime ledger indicates sufficient residual budget in the current decision window. Table 4 summarizes the SAL-aware scheduling policy and the minimum evidence fields required to replay decisions from logs.

Feasibility checks are computed from auditable quantities only: the queued airtime sum $\sum$ToA, the per-packet airtime estimate $\tau(m,p)$ from Section 5.1, and the next feasible transmit/receive opportunity under the active link profile (continuous for LAN/Wi-Fi; windowed for LPWAN such as LoRaWAN).

SAL transition triggers and hysteresis follow the policy defined in Section 5.1. Here we focus on the per-class scheduling rules and the minimal evidence required to replay queueing and transmission decisions under SAL-aware prioritization (Table 4).

To avoid assuming substantial endpoint storage, the authoritative audit trail is assembled at the service plane (gateway/server) from transmission metadata and endpoint acknowledgements, while endpoints retain only minimal local markers required for safety and eventual upload.

Table 4
SAL-aware scheduling policy and minimal audit evidence (summary).

Traffic class	SAL-2 policy	SAL-1 policy	Minimal evidence (assembled at service plane)
Time beacons (TIME_STATE)	Periodic; highest priority	Preserved; highest priority	endpoint_id/ group_id, ref_seq, T_ref tag, integrity_state, tx_ts, rx_ts/quality
MCAS control + completion (EXECUTE_AT/ ABORT/ CONFIRM)	Deadline-ordered; bounded retries; non-MCAS admitted only if slack	Reserved airtime; deadline-ordered; confirmations prioritized	cmd_id, class = MCAS, deadline (Δ_M/Δ_ACK), enqueue_ts, tx_profile, tx_attempts, ack/ confirm_ts, status_code
Non-MCAS traffic	Best-effort; admitted only with slack	Throttled/paused; buffered if needed	class = non-MCAS, payload_size, enqueue_ts, tx_ts (if any), drop/defer reason

5.3. Minimal evidence logs for audit

To make acceptance decisions reproducible, TCA[2] specifies a minimal evidence log schema aligned with Tables 1 and 2. The schema is intentionally distributed and size-bounded: resource-constrained endpoints keep only a small local ring buffer (e.g., in EEPROM), while the gateway/service node maintains the authoritative transaction timeline.

At minimum (Table 5), the collected logs SHALL allow an offline auditor to reconstruct: (i) reference-source selection and integrity-state changes, (ii) SAL transitions and their triggers, (iii) MCAS transactions (issue, transmit attempts, and confirmations) against declared deadlines Δ, and (iv) endpoint time-error samples $\widehat{e}(t_k)$-or explicit missing-sample markers under the sampling policy used for $A(\theta)$. Missing required evidence is treated conservatively as non-attainment unless independently corroborated (Section 3.2). Section 7 uses these records to compute acceptance metrics (A(θ), R(Δ)) and to summarize deployment evidence.

Illustrative SAL-2→SAL-1 episode. Suppose that, over W_{switch} = 120 s, the windowed airtime utilization U_W remains above U_{high} while the queue contains pending TIME_STATE frames, MCAS confirmations, and a larger backlog of non-MCAS telemetry. The scheduler projects that, if non-MCAS admission continues, some MCAS confirmations will exceed Δ_{ACK} = 30 s under the remaining airtime budget. The service node therefore records a SAL_TRANSITION event with fields such as {t_event, SAL_from = 2, SAL_to = 1, trigger_type = capacity, U_W, queue_depth, integrity_state = HEALTHY} and immediately pauses non-MCAS admission, reserving airtime for TIME_STATE and MCAS traffic. Recovery to SAL-2 requires $U_W < U_{low}$ for $W_{recover}$, together with MCAS attainment meeting the SAL-2 target over the same interval. This example is representative and is intended to show how auditable degradation is reconstructed from the retained evidence fields.

6. Implementation notes

This section summarizes the implementation details needed to reproduce the behaviors in Sections 3–5: integrity arbitration/holdover, SAL/MCAS enforcement, and evidence logging. We focus on interface semantics and the minimal log schema so alternative hardware or network stacks can be substituted without changing acceptance recomputation.

Some deployments further bound per-message occupancy by encoding MCAS messages in a fixed-length Compact State Vector (CSV) (Appendix A3). This encoding is non-normative and does not affect acceptance metrics.

Table 5
Minimal evidence log schema for offline acceptance recomputation and SAL explanation.

Record type	Produced by	Minimum required fields (required)	Used for
ARBITRATION_STATE (=REF_SELECT)	Service node/ gateway	t_event, selected_source, integrity_state, per_source_quality_flags, observed_offsets_summary, tau_int, window_id, [reason_code]	Explain integrity screening; reproduce source selection
SWITCH/HOLDOVER_EVENT	Service node	t_event, from_source, to_source (or HOLDOVER), reason_code, [supporting_evidence_ref]	Explain switches/holdover
SAL_TRANSITION	Service node	t_event, SAL_from, SAL_to, trigger_type (capacity/integrity), U_W, queue_depth, integrity_state	Explain SAL_state; reproduce degradation triggers
MCAS_TX	Gateway/server	msg_type, cmd_id, endpoint_id (or group_id), class = MCAS, Delta (Δ_M or Δ_ACK), issue_ts, enqueue_ts, tx_profile (or phy_profile_id), tx_attempts, ack/confirm_ts, status_code, [T_exec], [failure_cause]	Recompute R(Δ) (MCAS)
TIME_SAMPLE	Observer/ endpoint/audit tool	endpoint_id, t_sample (t_k), e_hat(t_sample) (or T_disp with T_ref), theta, sample_status (OK/MISSING/INVALID), observation_method, [uncertainty_bound], [observer_id]	Recompute A(θ)
AIRTIME_LEDGER	Gateway/server	t_window_start, T_obs, direction (DL/UL), gateway_id, link_profile_id (or phy_profile), sum_ToA, pkt_count	Recompute U_W; support capacity triggers
ENDPOINT_HEALTH (=ENDPOINT_STATE)	Endpoint/ gateway	endpoint_id, t, sync_state, holdover_flag, last_good_sync_time, [receiver_quality], [battery], [RSSI/SNR]	Diagnose; corroborate missing samples
CONFIG_SNAPSHOT (recommended)	TSN/service plane	schema_version, key thresholds (θ, Δ_M, Δ_ACK, τ_s), SAL parameters (U_high/U_low, T_obs, W_switch/W_recover), tau_int, active link_profile_id	Make recomputation reproducible

Note: Endpoints store only minimal local state (e.g., an SAL level indicator and a small MCAS status cache); the rest is recorded at the gateway/service node.

6.1. Components and responsibilities

TCA^2 is deployed as a service plane plus site gateways: the service plane ingests time sources, runs arbitration/holdover, and publishes the reference time and control schedules; site gateways terminate the LPWAN link and execute local scheduling and logging near the radio.

Time Service Node (TSN). The TSN maintains the system reference time $T_{\mathrm{ref}}(t)$ and an integrity state (e.g., HEALTHY/SUSPECT/HOLDOVER) via multi-source arbitration and explicit holdover logic. It exports (i) $T_{\mathrm{ref}}(t)$, (ii) the current integrity state and switch rationale, and (iii) the versioned configuration parameters that define acceptance thresholds and SAL transitions. The TSN also serves as the **timestamp authority** for audit logs and for measuring command issue/completion latencies, ensuring that $R(\Delta)$ is computed on a consistent timebase.

Site gateway/link server (packet-plane). This component performs packet routing and link-profile-aware scheduling for downlink/uplink traffic (Section 5); it is explicitly not a time authority. For each transmission attempt, the gateway/server records the active link profile and minimal diagnostics (e.g., LPWAN: SF/BW/CR, RSSI/SNR; LAN/Wi-Fi: rate/queueing/loss), payload length, attempts, and timestamps needed to recompute occupancy/utilization and to attribute deadline misses to queueing versus link failure.

Timecode Adapter. Given $T_{\mathrm{ref}}(t)$ and profile metadata (e.g., timecode profile, time-zone/DST policy, integrity flag), the adapter deterministically renders a national LF timecode representation (frame/symbol stream or baseband waveform) suitable for Path B dissemination.

Endpoints. Endpoints include (i) legacy LF clocks that passively receive tower LF (Path A) and/or site-local LF augmentation (Path B), and (ii) digital endpoints on LPWAN (Path C) that consume time beacons and MCAS commands and emit acknowledgements and health telemetry. To remain feasible on constrained hardware, endpoints are required only to maintain a minimal local state (sync/holdover/SAL flags) and to emit compact evidence (sequence numbers, status codes, and optional event timestamps); durable audit logs are stored on the service plane and/or gateways.

Observer/audit service. The observer aggregates telemetry and transaction logs to recompute $A(\theta)$ and $R(\Delta)$, assign fault classes, and reconstruct the rationale for integrity switches and SAL transitions. Optional external observations (e.g., privacy-preserving **vision-assisted** reading for analog clocks, using a replaceable observer module) are treated strictly as evidence channels, producing derived offsets and uncertainty bounds rather than becoming control-time authorities.

6.2. Interface conventions (minimal, implementation-agnostic)

To make the behavior in Sections 3–5 independently re-implementable, TCA^2 constrains only interface semantics and evidence timebases, not internal component structure. Every command carries a globally unique msg_id, a class label (MCAS vs non-MCAS), and an explicit deadline Δ; time-aligned actions additionally carry a reference execution time T_{exec}. Completion evidence is timestamped on the TSN timebase (send_ts for downlinks, recv_ts for uplinks), so that $R(\Delta)$ is computable without trusting endpoint clocks. Every time-error observation records the observation method (telemetry/manual/vision-assisted), the scheduled sampling index t_k (or window_id), and the missing-sample policy used when computing $A(\theta)$. Log events are schema-versioned and include the active PHY/profile identifier whenever airtime accounting is used.

6.3. Minimal audit logs aligned with the service contract

Acceptance is considered decidable only if the minimum evidence fields in Table 5 are recorded; sampling or throttling is allowed only when the acceptance metrics remain recomputable from retained logs. Logs are event-based, schema-versioned, and share a common timebase (TSN reference time for downlinks; TSN receive time for uplinks).

Missing required fields are treated conservatively as the corresponding criterion not satisfied for the affected metric interval, unless independently corroborated. To accommodate constrained endpoints, TCA^2 uses tiered logging: endpoints keep a small rolling buffer of local state and command outcomes, while the TSN/gateway persists the authoritative audit trail used for offline recomputation.

When stronger reproducibility is needed for adapter validation, immutable artifacts (e.g., frame_id or corpus_id, and optionally a hash) are generated and stored at the service/adapter side, not on leaf endpoints.

6.4. Compliance, safety, and privacy boundaries

TCA^2 is regulation-aware and intentionally non-prescriptive about physical transmission implementations. For LPWAN operation, deployments follow region-specific spectrum access rules and profiles; the scheduler therefore treats airtime and duty budget as first-class operational constraints rather than corner cases. For local LF augmentation (Path B; Fig. 4(a)), we specify only the functional interface and the compliance gate; this paper does not provide wiring, antenna, or injection procedures. Security boundaries are explicit: TCA^2 does not assume any single time source is trustworthy, and it treats anomalies and prolonged unavailability as integrity signals handled by arbitration and SAL

transitions.

Optional visual auditing (human- or AI-assisted observer). TCA^2 treats visual observation as an observation method, not as a control-time authority. Each retained visual observation must specify the observation method, the scheduled sample instant or window, the derived offset or coarse device-state cue, and an uncertainty bound compatible with Section 3.2(M1). Under this normalization, human and AI-assisted observations share the same evidentiary requirement: both must yield a derived observation plus an explicit uncertainty budget; otherwise they are retained only for maintenance triage. In this study, vision-assisted outputs are used for exploratory maintenance support and anomaly triage, not as a primary authority for computing R(Δ), and only as time-service evidence where a site-approved procedure and uncertainty declaration are available. Recognition reliability remains sensitive to viewpoint, pose, reflections, illumination, and motion, so the capture pipeline follows privacy-by-design and retains only derived offsets or bounded state observations under local policy.

6.5. Stratified resilience and fail-safe operation

Resilience. When upstream connectivity is lost (SAL-0), the TSN, gateways, and endpoints enter explicit holdover; endpoints may continue pre-scheduled actions while bounded-error claims are suspended, and the interval is marked as time-unassured in the audit trail. Higher availability can be achieved via an active–standby TSN deployment, while gateways remain replaceable and operational state is reconstructed from auditable logs and endpoint non-volatile markers. Optional local "break-glass" overrides (e.g., a physical switch or serial console) force a safe MCAS state and are recorded as operator interventions.

7. Evaluation

We evaluate TCA^2 with four evidence layers, each supporting a scoped claim: DS-1 verifies adapter portability via offline conformance/replay; DS-2 measures audited indoor synchronization for legacy clocks (tower LF vs site-local LF augmentation); DS-3 reports an anonymized footprint snapshot to ground evidence handling at scale (not a controlled benchmark); and DS-4 demonstrates packet-plane MCAS command evidence in a controlled testbed.

7.1. Datasets and methodology

Our evidence is explicitly layered rather than monolithic. No single dataset is intended to instantiate the entire TCA^2 stack, and each dataset is used only for the claim it can support. DS-1 addresses adapter semantic portability only. DS-2 uses legacy display-only clocks as a weakest-observability endpoint class and evaluates the time-service contract on one-way LF Paths A/B only. DS-3 is used only to ground field-scale evidence handling and operational heterogeneity, not to infer fleet-wide A(θ) or MCAS attainment. DS-4 instantiates Path C and provides the only direct evidence for deadline-audited minimal-control attainment R(Δ). The overall paper claim therefore follows from the composition of these scoped evidence layers rather than from a single end-to-end deployment. Table 6 summarizes the datasets, and Table 7 maps claims to evidence and scope boundaries.

Sections 7.2–7.5 report DS-1–DS-4 outcomes using the evidence definitions above; Section 7.6 summarizes integrity and degradation evidence hooks.

Table 6
Evaluation datasets and evidence outputs (DS-1/DS-2/DS-3/DS-4).

Dataset	Scope	Scale	Duration	Primary outputs
DS-1	Adapter conformance + replay	3 national LF profiles (BPC/DCF77/JJY)	24 h per profile + structured edge cases	bit/field-level diffs; replay verification; portability evidence
DS-2	Indoor field trial (single building)	30 endpoints (legacy clocks)	4 weeks	S_sync (audit-window pass rate); TTFS; energy/day
DS-3	Operational deployment snapshot	7848 terminals (22 provinces)	2020-2025 snapshot	deployment footprint; upgrade recency (last-upgrade year distribution); repeat-procurement concentration (rho_P, rho_T) from contract records
DS-4	Path C mechanism test (controlled) (same contract parameters; $\Delta_{ACK}=\Delta_M=$ 30 s)	LAN/Wi-Fi: hundreds of endpoint agents (software emulators; up to 784 observed); constrained LPWAN-profile: optional emulation mode	Complex network structure, NAT, VPN, laboratory tests distributed across three regions. Baseline 30 min + run 5 h	R_{ACK} (Δ_{ACK}), R_M (Δ_M), timeouts, and SAL transition evidence recomputed offline from schema-versioned JSONL logs

Note: DS-1 provides portability evidence via offline conformance (bit/field-level diffs against ground truth) and replay verification without over-the-air transmission. DS-2 reports S_{sync} as the audit-window pass rate; it serves as a practical proxy for $A(1s)$ on second-resolution legacy displays when continuous error traces are unavailable. DS-3 summarizes operational deployment scale; it complements controlled trials but does not substitute for population-level reliability statistics unless corresponding log completeness and maintenance aggregates are available.

7.2. Adapter portability via offline conformance and replay

DS-1 evaluates Timecode Adapter portability by verifying that country-specific LF timecodes can be decoded and re-encoded correctly without requiring any over-the-air transmission (Fig. 4(b)). The adapter is treated as a deterministic profile: given a system time tag and declared civil-time metadata (e.g., DST/leap indicators when applicable), it produces a 1-min timecode frame (symbol stream or baseband waveform) for the selected national scheme. In conformance mode, a minute-frame corpus is fed to the encoder under test, and the emitted frames are decoded by a separate reference decoder to recover fields; recovered fields are compared against the input annotations, producing a field/bit-level diff report when mismatches occur. The default corpus covers a 24-h sweep (1440 min frames) per timecode profile and is augmented with structured edge cases that exercise specification-defined transitions (e.g., DST changes) as well as robustness to missing or malformed markers.

In replay mode, the adapter output (symbol stream or baseband waveform) is fed into a receiver/decoder test chain (software or hardware front-end) to verify lock behavior and field recovery under the exact deployment profile, again without RF transmission.

We report DS-1 outcomes as (i) conformance pass/fail over valid vectors, (ii) mismatch counts and categories, and (iii) replay lock/recovery statistics, and we treat any mismatch on valid vectors as a portability defect for the corresponding adapter profile. To make DS-1 reproducible, each test run records the adapter profile identifier, adapter version, corpus identifier, and reference decoder version; large artifacts (e.g., emitted waveforms) are kept at the test harness when needed and are not assumed to be logged by resource-constrained endpoints.

Table 7
Claim-to-evidence mapping and scope boundaries.

Claim	Supported by	What it demonstrates	What it does not demonstrate	Why sufficient in context
National LF profiles can be normalized and re-rendered without changing contract semantics	DS-1	Adapter portability and semantic correctness	Indoor RF coverage or command attainment	Portability is an adapter property, not a propagation result
The time-service contract remains decidable on weakest-observability one-way legacy endpoints	DS-2	Audited time acceptance on display-only clocks over Paths A/B	R(Δ), continuous error traces, or universal coverage radius	DS-2 targets the acceptance interface, not bidirectional control
The evidence workflow can be operated at field scale across heterogeneous deployments	DS-3	Footprint, lifecycle heterogeneity, and evidence-handling plausibility	Fleet-wide A(θ) or MCAS statistics	DS-3 is an operational grounding layer, not a controlled benchmark
Packet-connected endpoints can provide auditable minimal-control attainment and SAL evidence	DS-4	R(Δ), ACK/CONFIRM evidence, late-confirm distinction, SAL events	Population-level field LoRaWAN performance or LF-only control	Minimal-control attainment requires packet-plane completion evidence
TCA^2 constitutes a hybrid assurance architecture under degraded hybrid links	DS-1 + DS-2 + DS-3 + DS-4	Unified contract across portability, weakest-interface time acceptance, field-scale workflow grounding, and packet-plane minimal control	High-dynamic hard real-time continuous autonomy	The claim is compositional and evidence-bounded by design

DS-1 establishes adapter semantic portability and implementation correctness; RF coverage and indoor propagation effects are evaluated separately in DS-2/DS-3.

7.3. Indoor field trial on weakest-observability legacy endpoints

DS-2 intentionally selects a weakest-observability endpoint class: legacy display-only LF clocks on one-way Paths A/B. The purpose is not to represent the full endpoint universe of TCA^2, but to test whether the time-service contract remains decidable when the acceptance interface offers only second-quantized scheduled observations and no uplink telemetry. We evaluate this setting in a four-week single-building field trial covering N = 30 legacy clocks deployed across corridors and rooms, comparing tower-only reception (Path A) with tower plus site-local LF augmentation (Path A + B). Because these endpoints do not emit continuous offset telemetry, we report an audit-time compliance metric rather than continuous-time error distributions. This audited estimator is the natural instantiation of A(θ) under a display-only evidence interface, and it does not attempt to infer between-audit drift.

At each scheduled audit instant t_a, we read the displayed time $T_{\mathrm{disp},i}(t_a)$ and form an error sample $e_i(t_a) = T_{\mathrm{disp},i}(t_a) - T_{\mathrm{ref}}(t_a)$, where T_{ref} is the system reference time. For DS-2 legacy clock endpoints, the error sample at each audit instant is second-quantized (integer-second display ticks). To remain conservative under second-quantized readouts, an endpoint is counted as synchronized only if $|e_i(t_a)| + u_{e,i} \le \theta$ with $\theta = 1$ s, where $u_{e,i}$ captures audit/readout uncertainty. We define the audited synchronization availability $S_{\mathrm{sync}}(t_a)$ as the fraction of endpoints satisfying the above rule at t_a; missing scheduled audits are treated as unsynchronized ("missing = unavailable").

Fig. 6 summarizes the weekly audited synchronization availability. Under tower-only reception, $S_{\mathrm{sync}} \approx 0.70$ ($\approx 21/30$) due to persistent indoor blind spots. With site-local LF augmentation, S_{sync} rises to 0.933 (28/30) in W2 and reaches 1.00 (30/30) in W3–W4, indicating that local augmentation closes the site's last-meter reception gaps under the audited criterion.

We additionally report time-to-first-synchronization (TTFS) for cold starts, defined as the elapsed time from power-on until an endpoint decodes a valid LF minute frame and completes its first adjustment. Across observed cold starts, TTFS ranges from 2 to 10 min, consistent with reception conditions and minute-level framing/acquisition latency.

To support operational planning, we measure endpoint energy consumption under the DS-2 workload and report a daily energy budget of 14.83–16.50 mWh/day. With a voltage-windowed useable energy of 32.244–34.940 Wh above a 2.2 V cutoff (Table A6.3), this budget implies a projected service life of 5.34–6.45 years, i.e., 15.5–18.7 battery swaps per 100 device·year.

Fig. 7 reports these planning quantities as derived signals, so that deployments can re-compute the same metrics under different battery capacities or cutoff policies without changing the acceptance definitions. We emphasize that S_{sync} is an audit-time acceptance measure for

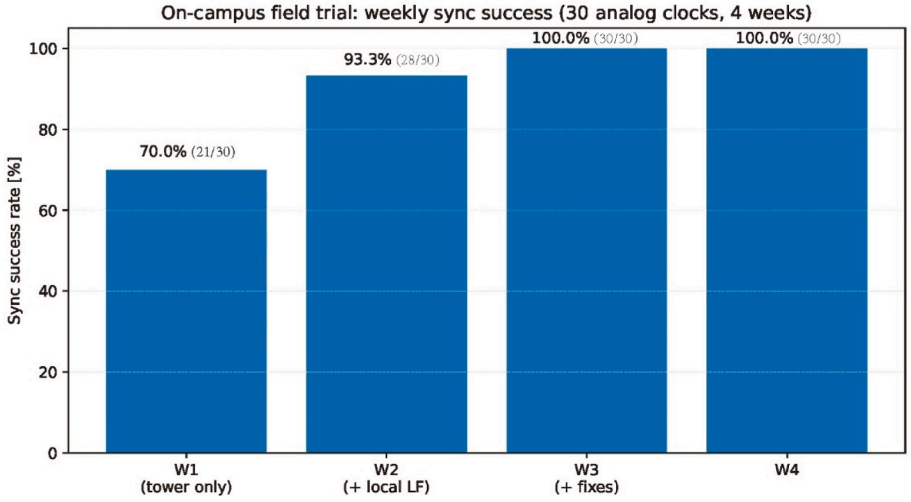

Fig. 6. Audited synchronization availability S_{sync} in DS-2 (30 legacy display-only clocks, 4 weeks). An endpoint is counted as synchronized at an audit instant if $|e_i| + u_{e,i} \le 1$ s; missing scheduled audits are treated as unsynchronized.

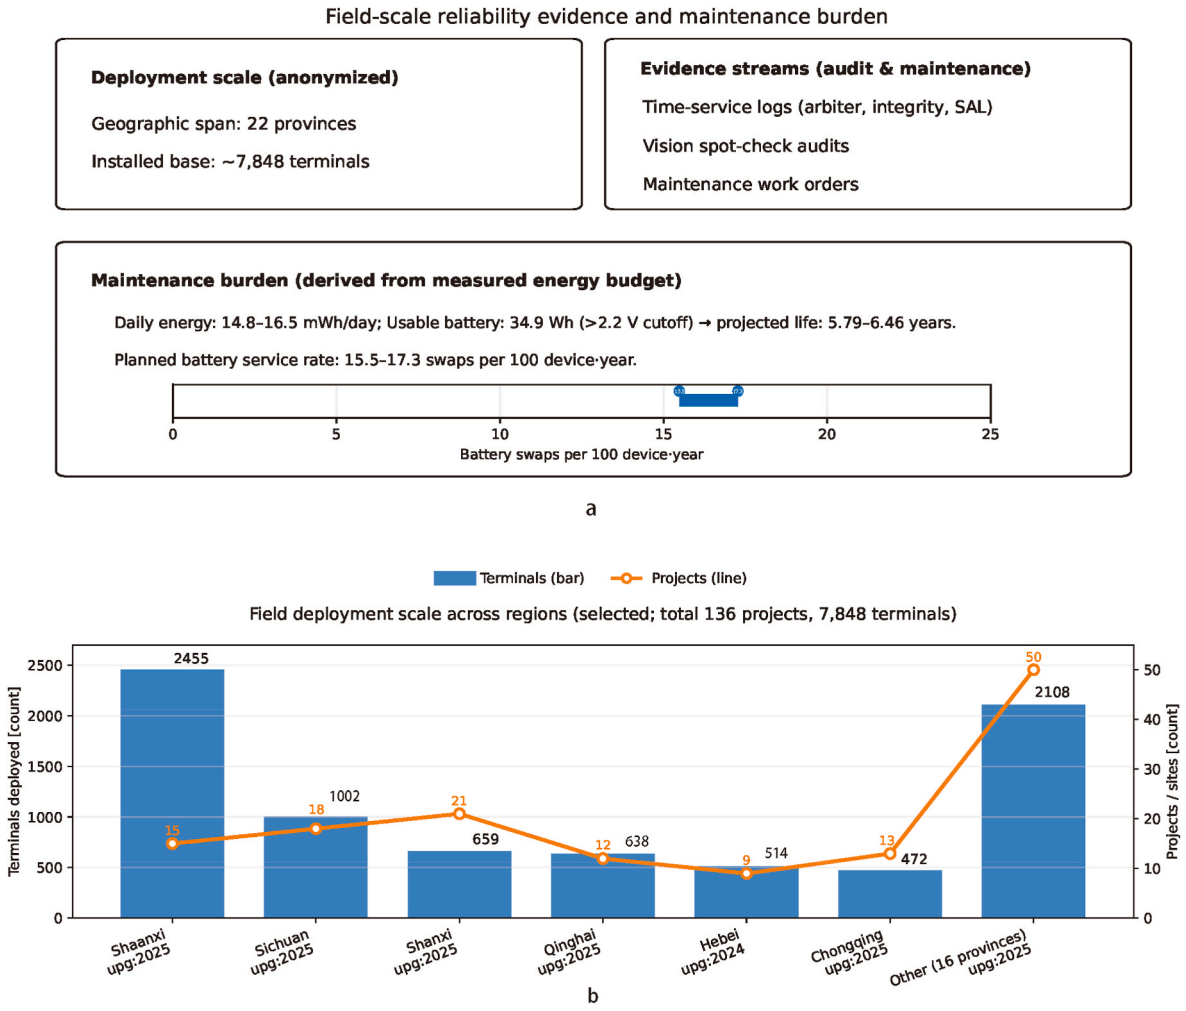

Fig. 7. Planning signals and footprint evidence. *(a) DS-2 energy measurements imply a daily budget of 14.83–16.*50 mWh*/day and a projected lifetime of 5.34–6.45 years given 32.244–34.*940 Wh *useable energy (2.2 V cutoff;* Appendix Table A6.3*). (b) DS-3 anonymized footprint visualization (7848 terminals over 136 projects; region groups shown). Panels (a)–(b) are different evidence layers and should not be interpreted as one homogeneous full-stack deployment.*

non-instrumented, second-quantized displays; instrumented digital endpoints can instead report continuous offset samples to compute $A(\theta)$ directly under a declared sampling interval.

7.4. Large-scale deployment footprint

DS-3 is an anonymized, registry-based operational snapshot drawn from administrative evidence (deployment registry, procurement/contract records, and upgrade/maintenance records). It is used to ground the evidence workflow at field scale rather than to provide a controlled performance study. The snapshot covers $N_P = 136$ projects (schools/exam sites) across 22 provincial-level regions, totaling $N_T = 7848$ terminals, with $N_C = 214$ procurement contracts recorded up to 2025.

For auditability and privacy, we retain only non-personal fields per project: region, terminal count, first commissioning year, and the most recent recorded upgrade year; earlier prototype deployments with incomplete fields are excluded. DS-3 is therefore used to characterize footprint and lifecycle activity from administrative evidence, and it is not used to infer fleet-wide A(θ) or MCAS attainment without uniform endpoint logs.

To avoid treating DS-3 as a scale-only anecdote, we report registry-derived indicators that are exactly recomputable from the retained fields.

(i) Repeat-procurement concentration. Let N_{multi} be the number of projects with at least two contracts within the statistics window, and let T_{multi} be the terminal count within those projects. Define $\rho_p = N_{multi}/N_p$ and $\rho_t = T_{multi}/N_t$. In DS-3, $\rho_p = 39/136 = 0.287$ and $\rho_t = 4991/7848 = 0.636$. Repeat-procurement projects average 128 terminals/project versus 29 for single-procurement projects (a 4.34 × density gap).

(ii) Upgrade currency. The most recent recorded upgrade year in the registry is 2024–2025; at the region-level snapshot visualized in Fig. 7(b), 334/7848 terminals (93.4%) belong to regions whose latest recorded upgrade is 2025, while 514/7848 (6.6%) belong to a region whose latest recorded upgrade is 2024. This is a coarse lifecycle indicator and is not interpreted as a timing-accuracy metric.

(iii) Staged deployment modes (lower bound). Procurement line-item flags indicate limited shipment of auxiliary components. Aggregated at project level, 27/136 projects (covering 1442/7848 terminals) record at least one auxiliary item (site-local server, LoRa time-sync/relay nodes, or Wi-Fi time-synchronization modules), consistent with incremental rollouts; DS-3 does not include uniform per-device logs for these subsets and thus does not replace DS-2/DS-4.

DS-3 therefore complements DS-2 by exposing operational heterogeneity and validating that the evidence workflow can be operated at

field scale, without being interpreted as a controlled benchmark of A(θ) or MCAS attainment.

7.5. DS-4 controlled packet-plane testbed (LAN/Wi-Fi)

DS-4 instantiates Path C as a controlled packet-plane testbed on an IP LAN/Wi-Fi, where bidirectional evidence (ACK/CONFIRM) is available. A constrained LPWAN-profile interface (labeled "lora" in the logs) is included only as an emulation mode to exercise capacity-driven degradation and SAL evidence. We run a gateway implementation together with a pool of lightweight endpoint agents (software emulators). Each MCAS command follows the minimal transaction pattern: an EXECUTE_AT issue is followed by an ACK and a CONFIRM, while background uplink frames exercise the observation layer. The contract parameters recorded in RUN_HEADER are Δ_ACK = 30 s and Δ_M = 30 s; all MCAS events (MCAS_ISSUE/ACK/CONFIRM, *_TIMEOUT, and SAL_TRANSITION) are timestamped on the gateway timebase and stored as JSONL records conforming to the minimal evidence schema (Table 5).

The dataset consists of a 30-min baseline followed by a 5h run under the same deadlines. Command injection rates are 2.0 commands/min on the LAN/Wi-Fi profile and 0.5 commands/min on a constrained LPWAN-profile interface (emulation mode with windowed receive and airtime-ledger budgeting). During the run, the gateway tracks hundreds of endpoint agents on the LAN/Wi-Fi profile (up to 784 agents observed in RUN_HEARTBEAT); the constrained profile is included to exercise capacity-driven degradation and SAL evidence, not to benchmark field LoRaWAN performance.

We recompute control attainment post hoc using Section 3.2 (M3) definition under the conservative rule "missing evidence = not attained". We report $R_{\mathrm{ACK}}(\Delta_{\mathrm{ACK}})$ as the fraction of issued MCAS commands whose ACK evidence arrives within Δ_{ACK}, and $R_{\mathrm{M}}(\Delta_{\mathrm{M}})$ as the fraction whose CONFIRM evidence is deadline-bounded and therefore counted (MCAS_CONFIRM.counted). For diagnostic clarity, we also report R_{exec} as the fraction of issued commands that were executed regardless of deadline (MCAS_CONFIRM.attained). On the LAN/Wi-Fi profile, baseline yields $R_{\mathrm{ACK}} = 1.000$ and $R_{\mathrm{M}} = 0.948$ ($R_{\mathrm{exec}} = 0.966$); the 5h run yields $R_{\mathrm{ACK}} = 0.997$ and $R_{\mathrm{M}} = 0.965$ ($R_{\mathrm{exec}} = 0.976$) over 593 issued commands. Late confirmations (attained = true but counted = false) are explicitly logged, illustrating that TCA^2 distinguishes execution from deadline-bounded auditable acceptance.

By the SAL thresholds in Table 1, this stress configuration corresponds to SAL-1 for MCAS control, while acknowledgements remain in the ≥0.99 range. The preceding commissioning baseline is used to populate the endpoint registry and validate end-to-end logging; inclusive baseline counts (including admission transients) are reported in the Supplementary materials.

For transparency and reviewer checkability, the supplementary material includes the raw DS-4 JSONL logs (baseline and run) and a dependency-free recomputation script that reproduces the exact counts and ratios above and outputs per-profile CSV/JSON summaries. DS-4 should be read as controlled evidence for packet-plane minimal-control attainment and SAL reconstruction, not as a population-level field LoRaWAN benchmark and not as evidence for deadline-audited control on LF-only legacy paths.

7.6. Integrity and degradation: evidence design and reproducibility hooks

To make integrity-driven behavior reviewer-checkable, TCA^2 logs integrity changes as schema-versioned events (see the minimal log schema in Table 5). Each episode forms a replayable chain detection → arbitration decision → holdover entry/exit (if any) → recovery-so switch rationale can be reconstructed offline.

DS-2/DS-3 cover the time service on one-way legacy endpoints (Paths A/B) and therefore cannot evaluate command attainment, since such clocks provide no ACK/CONFIRM evidence. MCAS command attainment $R(\Delta)$ is demonstrated only in DS-4 on the packet plane. Fleet-wide integrity analytics would require uniform export of arbitration/scheduling logs and time-error observations across sites; we treat this as a current limitation (Section 8).

8. Discussion and threats to validity

Measurement uncertainty. Acceptance is computed at the endpoint acceptance interface, so claims are bounded by the sampling policy and observation-channel uncertainty. In DS-2, legacy clocks are one-way and expose only second-quantized displays (no uplink telemetry or device logs); therefore we report an audit-time compliance metric and treat missing scheduled audits as unavailable. For indirect observation (human/vision), quantization, readout ambiguity, and audit-time alignment are bounded by an explicit u_e, and compliance is declared only under the conservative rule $|e| + u_e \leq \theta$.

This external observation does not indicate weak internal synchronization: the arbitration core and dissemination scheduler remain auditable via service-plane evidence (reference selection states and transmission schedules), while observation is used only to assess the realized physical behavior of display-only or mechanically actuated endpoints, which can deviate due to mechanism drift or installation bias. Vision remains an optional evidence channel and is not used as a primary authority in this study.

Generalizability. Indoor LF reception and any in-building augmentation are strongly site-dependent, influenced by building materials, geometry, EMI, and transmitter placement. Therefore, DS-2 demonstrates feasibility and a within-site uplift, but we do not claim a universal coverage radius or a portable "coverage number." TCA^2 addresses this by treating local LF augmentation as a compliance-gated module and by judging each site against the same acceptance rules $(A(\theta), R(\Delta), \mathrm{SAL})$ rather than against a single fixed propagation model.

Evidence scope and security boundaries. DS-3 provides operational scale indicators, but not the same granularity as a controlled trial, and we do not use it to make fleet-wide claims about full error distributions. Closing the evidence loop for integrity switching and congestion-driven degradation would require exportable, privacy-preserving arbitration and scheduling logs at consistent schemas across deployments, together with replayable perturbation traces; we treat this as a limitation of the current evidence set rather than as a requirement for the contract definitions themselves. The integrity-driven arbiter is designed for anomaly detectability and explainable switching under non-adversarial faults; we do not claim robustness against a powerful active **attacker** beyond what is explicitly evaluated. Likewise, log integrity and access control are deployment responsibilities: if required evidence is missing or tampered with, acceptance is weakened by design.

At the evidence-management layer, TCA^2 could be complemented by cryptographic and verifiable data-management mechanisms for privacy-preserving audit retrieval, cross-site evidence sharing, private set computation, and evidence-lifecycle verification. Recent work on blockchain-assisted verifiable encrypted search, federated encrypted data stores, reliable private set computation in decentralized storage, and verifiable machine unlearning provides relevant building blocks for such future extensions [26–29]. These mechanisms are complementary to the present work and are not evaluated as part of the transport-level time-control assurance claims.

Applicability boundary. The present evidence supports low-rate, deadline-bounded critical continuity under degraded hybrid links: time continuity, minimal guaranteed triggers, safe-state transitions, and auditable degradation. It does not yet constitute a validation of high-frequency, hard real-time, continuously closed-loop autonomy, nor does it establish deadline-audited control attainment for one-way LF-only legacy endpoints. In other words, LF dissemination in this paper is best understood as a resilient backbone for time continuity or coarse triggers, while packet-connected Path C provides the auditable control evidence.

9. Conclusion

This paper presents TCA^2 as a minimum-assurance time–control framework for GNSS-challenged indoor regions under degraded hybrid links. The principal contribution is not a new standalone timing carrier or LPWAN primitive, but an evidence-bound assurance contract that unifies multi-source time arbitration, minimum-control preservation, auditable degradation, and offline-recomputable acceptance. Time acceptance is expressed as $A(\theta)$ at the endpoint interface; minimal-control acceptance is expressed as $R(\Delta)$ for packet-connected endpoints within MCAS.

TCA^2 separates roles across dissemination paths. One-way LF paths (A/B) provide resilient time continuity for legacy or weakest-observability endpoints, while Path C provides the ACK/CONFIRM evidence required for deadline-audited minimal control. Layered validation supports this claim compositionally: DS-1 establishes adapter portability, DS-2 demonstrates auditable time acceptance on legacy display-only endpoints, DS-3 grounds the evidence workflow at field scale, and DS-4 demonstrates packet-plane MCAS attainment and SAL evidence.

Accordingly, the present evidence supports low-rate critical continuity, safe-state triggers, and explainable degradation under constrained conditions; it does not claim high-dynamic hard real-time autonomy or deadline-audited control on LF-only legacy paths. This boundary is intentional: TCA^2 is designed to keep a minimum service contract decidable when evidence becomes incomplete, not to mask degraded conditions behind best-effort behavior.

CRediT authorship contribution statement

Hui Dengfeng: Writing – review & editing, Writing – original draft, Validation, Software, Resources, Project administration, Methodology, Formal analysis, Conceptualization. **Shamala K. Subramaniam:** Supervision, Resources, Project administration, Conceptualization. **Lili Nurliyana binti Abdullah:** Writing – review & editing, Supervision, Formal analysis, Conceptualization. **Abdullah bin Muhammed:** Supervision, Resources, Conceptualization. **Liu Hongyan:** Validation, Software, Project administration, Funding acquisition. **Wang Miao:** Validation, Resources, Investigation. **Wang Zijun:** Validation, Investigation, Data curation. **Hui Yuchun:** Visualization, Software, Formal analysis, Data curation. **Zhu Junyu:** Visualization, Validation, Investigation. **Liu Zhongliang:** Writing – review & editing, Supervision, Resources, Methodology.

Ethics statement

This study involves no human-subject interventions. When visual auditing is used, capture is restricted to the clock face and only derived timing offsets (with uncertainty bounds) are retained.

Declaration of generative AI and AI-assisted technologies

During the preparation of this work, the authors used OpenAI ChatGPT to assist with language editing, grammar checking, and improving readability. The authors reviewed and edited the generated text as needed and take full responsibility for the content of the manuscript.

Funding statement

This research received no external competitive grant funding. Beijing Wholeway Centron Technology Co., Ltd. provided non-financial support in the form of equipment, operational access, administrative/data support, and authorization to use relevant intellectual property. The authors are responsible for the study design, analysis, interpretation, and manuscript conclusions.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Hui Dengfeng reports a relationship with Beijing Wholeway Centron Technology Co., Ltd. that includes: board membership, employment, and equity or stocks. Liu Hongyan reports a relationship with Beijing Wholeway Centron Technology Co., Ltd. that includes: board membership, employment, and equity or stocks. Hui Yuchun has patent #Software copyright registration 2023SR0066598 — Satellite Clock Signal Enhancer (Repeater) System, V1.0 licensed to Beijing Wholeway Centron Technology Co., Ltd. Hui Yuchun has patent #Software copyright registration 2023SR0066596 — Satellite-Clock Radio-Wave Signal Enhancer (Wi-Fi Version) System, V1.0 licensed to Beijing Wholeway Centron Technology Co., Ltd. Hui Dengfeng has patent #Utility model patent ZL202021033137.4 (published as CN211956148U) — An Examination-Room Clock Time-Synchronization Device issued to Beijing Wholeway Centron Technology Co., Ltd. Hui Dengfeng has patent #Software copyright registration 2023SR1509016 — Hongwei Xianchuang WBC-LifeTime Integrated Satellite Radio-Wave Clock Time-Synchronization System, V1.0 issued to Beijing Wholeway Centron Technology Co., Ltd. Hui Dengfeng has patent #Software copyright registration 2017SR241233 — Wholeway Centron Standard Time Signal Compensation and Synchronization System, V1.0 issued to Beijing Wholeway Centron Technology Co., Ltd. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.array.2026.100866.

Data availability

Data will be made available on request.

References

[1] Psiaki ML, Humphreys TE. GNSS spoofing and detection. Proc IEEE 2016;104(6):1258–70. https://doi.org/10.1109/JPROC.2016.2526658.
[2] Meng L, Yang L, Yang W, Zhang L. A survey of GNSS spoofing and anti-spoofing technology. Remote Sens 2022;14(19):4826. https://doi.org/10.3390/rs14194826.
[3] Joint Committee for Guides in Metrology (JCGM). International vocabulary of metrology—Basic and general concepts and associated terms (VIM). JCGM 2012;200:2012. https://doi.org/10.59161/JCGM200-2012.
[4] Joint Committee for Guides in Metrology (JCGM). Evaluation of measurement data—Guide to the expression of uncertainty in measurement (GUM). JCGM 2008;100:2008(E. https://doi.org/10.59161/JCGM100-2008E.
[5] Lombardi MA, Nelson GK. WWVB: a half century of delivering accurate frequency and time by radio. J Res Nat Inst Stand Technol 2014;119:25–54. https://doi.org/10.6028/jres.119.004.
[6] Bauch A, Hetzel P, Piester D. Time and frequency dissemination with DCF77: from 1959 to 2009 and beyond. PTB Mitt 2009;119(3):3–26.
[7] National Institute of Information and Communications Technology (NICT). "JJY standard time and frequency signal: transmission information," online resource. Available: https://jjy.nict.go.jp/jjy/trans/index-e.html. [Accessed 29 January 2026].
[8] National Time Service Center, Chinese Academy of Sciences (NTSC). "Facilities: BPC low-frequency time code time service station," online resource. Available: https://english.ntsc.cas.cn/research/facilities/. [Accessed 29 January 2026].
[9] ITU-R. Time and frequency services, including the use of GNSS. Geneva: International Telecommunication Union; 2021. Report ITU-R TF.2487-0.
[10] LoRa Alliance. "LoRaWAN® Specification v1.1," LoRa alliance. Available: https://resources.lora-alliance.org/technical-specifications/lorawan-specification-v1-1. [Accessed 29 January 2026].
[11] LoRa Alliance. "LoRaWAN® Regional Parameters RP002-1.0.5," LoRa alliance, Oct. 2025. Available: https://resources.lora-alliance.org/document/rp002-1-0-5-lorawan-regional-parameters. 2026-01-29.
[12] LoRa Alliance. Application layer clock synchronization specification TS003-2.0.0. LoRa alliance. Apr. 2022. Available: https://resources.lora-alliance.org/technical-s

pecifications/ts003-2-0-0-application-layer-clock-synchronization. [Accessed 29 January 2026].
[13] ETSI. Short range devices (SRD) in the frequency range 25 MHz to 1 000 MHz; part 2: harmonised standard for access to radio spectrum for non specific radio equipment. ETSI EN 300 220-2 V3.3.1 (2025-03), https://www.etsi.org/deliver/etsi_en/300200_300299/30022002/03.03.01_60/en_30022002v030301p.pdf; Mar. 2025. 2026-01-29.
[14] Adelantado F, Vilajosana X, Tuset-Peiró P, Martinez B, Melià-Seguí J, Watteyne T. Understanding the limits of LoRaWAN. IEEE Commun Mag 2017;55(9):34–40. https://doi.org/10.1109/MCOM.2017.1600613.
[15] Bor MC, Roedig U, Voigt T, Alonso JM. Do LoRa low-power wide-area networks scale?. In: Proceedings of ACM MSWiM; 2016. p. 59–67. https://doi.org/10.1145/2988287.2989163.
[16] Mills DL, Martin J. Network time protocol version 4: Protocol and algorithms specification. In: Burbank J, Kasch W, editors. Rfc 5905. IETF; Jun. 2010. https://doi.org/10.17487/RFC5905.
[17] Franke D, Sibold D, Teichel K, Dansarie M, Sundblad R. Network time security for the network time protocol. In: Rfc 8915. IETF; Sep. 2020. https://doi.org/10.17487/RFC8915.
[18] IEEE Instrumentation and Measurement Society. IEEE standard for a precision clock synchronization protocol for networked measurement and control systems. IEEE Std; 2020. https://doi.org/10.1109/IEEESTD.2020.9120376.
[19] IEEE. IEEE standard for local and metropolitan area networks—timing and synchronization for time-sensitive applications. In: IEEE std 802.1AS-2020; 2020. p. 1–421. https://doi.org/10.1109/IEEESTD.2020.9121845.
[20] Balakrishnan K, Dhanalakshmi R, Sinha BB, Gopalakrishnan R. Clock synchronization in industrial internet of things and potential works in precision time protocol: review, challenges and future directions. International Journal of Cognitive Computing in Engineering 2023;4:205–19. https://doi.org/10.1016/j.ijcce.2023.06.001.
[21] Joint Committee for Guides in Metrology (JCGM). Evaluation of measurement data—The role of measurement uncertainty in conformity assessment. JCGM 2012; 106:2012. https://doi.org/10.59161/JCGM106-2012.
[22] DeMedeiros K, Hendawi A, Alvarez M. A survey of AI-based anomaly detection in IoT and sensor networks. Sensors 2023;23(3):1352. https://doi.org/10.3390/s23031352.
[23] Yang C., Xie W., Zisserman A. It's about time: analog clock reading in the wild. Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR); 2022. p. 2498–2507. doi:10.1109/CVPR52688.2022.00254.
[24] Tu C, Chen M, Zhang L, Zhao L, Wu D, Yue Z. Towards efficient multi-granular anomaly detection in distributed systems. Array 2024;21:100330. https://doi.org/10.1016/j.array.2023.100330.
[25] Blais A, Couellan N, Munin E. A novel image representation of GNSS correlation for deep learning multipath detection. Array 2022;14:100167. https://doi.org/10.1016/j.array.2022.100167.
[26] Guo Y, Zhang C, Wang C, Jia X. Towards public verifiable and forward-privacy encrypted search by using blockchain. IEEE Transactions on Dependable and Secure Computing 2023;20(3):2111–26. https://doi.org/10.1109/TDSC.2022.3173291.
[27] Guo Y, Xi Y, Wang H, Wang M, Wang C, Jia X. FedEDB: Building a federated and encrypted data store via consortium blockchains. IEEE Transactions on Knowledge and Data Engineering 2024;36(11):6210–24. https://doi.org/10.1109/TKDE.2023.3341149.
[28] Y. Guo, Y. Xi, Y. Zhang, M. Wang, S. Wang, and X. Jia, "Toward efficient and reliable private set computation in decentralized storage," IEEE Transactions on Services Computing, vol. 17, no. 5, pp. 2945–2958, 2024. doi:10.1109/TSC.2024.3441471.
[29] Y. Guo, Y. Zhao, S. Hou, C. Wang, and X. Jia, "Verifying in the dark: Verifiable machine unlearning by using invisible backdoor triggers," IEEE Transactions on Information Forensics and Security, vol. 19, pp. 708–721, 2024. doi:10.1109/TIFS.2023.3328269.