

AI Quality Monitor

Doc type: technical · **Version:** v0.1 · **Status:** published · **Module slug:** ai-quality-monitor
Exported: 2026-05-15 11:13 UTC · **By:** anonymous

AI Quality Monitor – Technical Specification

1. Module Purpose & Scope (Authoritative)

AI Quality Monitor is a governed quality-assurance, documentation, and clinical decision-support system that continuously evaluates how clinical and operational work is performed and produces review-first outputs to improve patient safety, documentation completeness, and compliance readiness. It operates in the background — detecting missed steps, identifying documentation gaps, and generating auditable draft artefacts aligned to appointments and workflows — without acting as a surveillance or performance-management tool. It sits at the intersection of the clinical encounter and the practice's governance obligations, providing structured, explainable outputs that always require human review before they have effect.

It governs:

- Zone-based quality monitoring (configuration, signal ingestion, derived-signal generation, exception findings)
- Review-first clinical documentation and decision-support drafts (clinical summaries, treatment plans, care plans, hygiene plans, patient-safe summaries)
- Governed evidence capture and retention linked exclusively to quality findings
- Remediation workflow creation and routing via Task Manager and Communication Hub

It explicitly does not:

- Monitor or score individual staff performance (prohibited — see §3.2)
- Make autonomous clinical decisions or replace clinical judgement (owned by the clinician; surfaced only through review-first outputs)
- Manage task lifecycle or notification delivery (owned by Task Manager and Communication Hub respectively)
- Govern access roles or permissions (owned by Access Manager)
- Own the patient record or finalise clinical notes (owned by the PMS / clinical records surface)

2. Ownership & Responsibilities

2.1 AI Quality Monitor IS Responsible For

- Zone-based monitoring configuration: defining zones, sensors, and workflow-alignment rules per zone
- Signal ingestion and transformation into derived operational signals (not raw input storage)
- Generating exception findings with explainability references and context linkage
- Governing evidence capture: clipping, retention windows, expiry enforcement, and RBAC-gated access
- Producing all review-first clinician-facing and patient-facing draft output types defined in §7

- Creating structured remediation artefacts (tasks, alerts) and routing them outbound
- Logging all AI suggestions, findings, and attribution events as immutable audit records
- Emitting quality indicators and zone-summary telemetry to dashboards and AI Guardian

2.2 AI Quality Monitor IS NOT Responsible For

- Live staff observation or supervision (prohibited — no module owns this)
- Generating disciplinary or HR evidence by default (out of scope; any governance extension requires a logged governance decision)
- Task lifecycle management — owned by Task Manager
- Notification delivery and escalation threading — owned by Communication Hub
- Role and permission management — owned by Access Manager
- Patient record storage or finalisation of clinical notes — owned by the PMS integration boundary
- Higher-level operational awareness aggregation — owned by AI Guardian
- Scoring or ranking individual practitioners — prohibited (see §3.2 and §7 AI Boundaries)

3. Core Objects (Normative)

3.1 Quality Finding (Canonical Artefact)

A Quality Finding is a governed digital artefact representing a single detected exception or documentation gap, generated from derived signals, with explainability references and context linkage to the originating appointment, call, sterilisation cycle, or rota window.

Minimum required fields:

- FindingID (UUID)
- ZoneID (FK to configured zone)
- AppointmentID / ContextID (FK to linked operational context)
- FindingType (protocol gap | documentation gap | timing anomaly | cross-verification failure | attribution event)
- DerivedSignalRef (reference to the derived signal that triggered the finding)
- ExplainabilityNote (human-readable statement of what was detected and why)
- Attribution (deterministic actor reference, or `Unattributed` where confidence is insufficient)
- FindingState
- CreatedAt
- EvidenceClipRef (nullable FK — only populated under §5.1)
- AuditTrail (immutable)

3.2 Quality Finding State Machine (Authoritative)

States:

- `Generated` — finding created by the system; no human action yet taken
- `UnderReview` — a governance-authorized user has opened the finding

- `Remediated` — a remediation task has been created and linked
- `Resolved` — finding closed following human confirmation
- `Dismissed` — finding closed without remediation action, with a required reason recorded
- `Expired` — finding and any attached evidence have passed the retention window and been purged

Rules:

- State transitions are auditable and time-stamped with actor identity
- A finding cannot return to `Generated` once it has moved to `UnderReview`
- `Dismissed` requires a mandatory reason field; dismissal is logged as an audit event
- `Expired` is system-enforced and cannot be reversed; evidence clips expire on the same schedule

3.3 Draft Output Artefact (Canonical Artefact)

A Draft Output Artefact is a review-first document generated by AI Quality Monitor and presented to an authorised human for review, editing, and approval before any effect is applied to the patient record, treatment plan, or patient-facing surface.

Minimum required fields:

- `DraftID` (UUID)
- `DraftType` (see §7 output types)
- `AppointmentID` (FK)
- `GeneratedBy` (AI model reference)
- `DraftState`
- `ProposedContent` (structured text)
- `ReviewedBy` (user/role — nullable until reviewed)
- `ReviewedAt` (nullable until reviewed)
- `Disposition` (`Accepted` | `Edited` | `Rejected` — nullable until reviewed)
- `AuditTrail` (immutable)

3.4 Draft Output Artefact State Machine (Authoritative)

States:

- `Pending` — generated; awaiting clinician review
- `InReview` — opened by the authorised reviewer
- `Accepted` — approved as-is and committed to the target surface
- `Edited` — modified by the reviewer and committed
- `Rejected` — declined by the reviewer; no content committed
- `Expired` — not actioned within the retention window; purged

Rules:

- A Draft Output Artefact **MUST NOT** be auto-finalised or auto-committed under any circumstance
- Transition to `Accepted` or `Edited` requires an explicit human action by an authorised role
- All transitions are logged with actor, timestamp, and (where `Edited`) a diff reference

- `Rejected` and `Expired` drafts are retained in the audit log but their content is not surfaced to downstream systems

4. Zone-Based Monitoring & Signal Model

4.1 Zone Configuration (Authoritative)

The module **MUST**:

- Require explicit per-zone enablement before any capture or monitoring is active in that zone
- Document each enabled zone in staff policy and patient signage records (governance gate)
- Enforce RBAC on zone configuration changes; all changes are audited

The module **MAY**:

- Support the following zone types where explicitly enabled: clinical (surgery rooms), reception / front-of-house, decontamination, back-office / admin, arrival / entry, meeting / conference room
- Allow zone-level configuration of which input signal types are active (audio, camera events, access feeds, system events)

The module **MUST NOT**:

- Activate any capture surface without explicit enablement and a corresponding audit record
- Operate hidden or unreviewable capture

4.2 Signal Ingestion & Derived Signals (Authoritative)

Raw inputs are transformed into derived operational signals. Derived signals — not raw inputs — drive findings.

Supported input types (where enabled per zone):

- Ambient audio (UniFi Protect preferred; USB microphone fallback)
- Camera event feeds (not continuous video viewing)
- Access / badge / attendance feeds; ANPR where lawful and consented
- Decontamination equipment telemetry APIs (where available)
- System events: appointments, tasks, forms, aftercare, care plans
- AI Concierge automated-action events (forms sent, tasks created, recovery outcomes from voice/call workflows — recorded as system-event signals and logged as AI Quality Monitor attribution events)
- AI Meeting Notes events (transcription completion, speaker attribution confidence scores, action-item extraction outcomes — recorded as system-event signals from the meeting / conference room zone)

Derived signal types include:

- Protocol step present / missing
- Documentation completeness flags
- Timing and sequencing anomalies
- Cross-verification results (Forms, Aftercare, Tasks, Care Plans, PMS)

The module **MUST**:

- Ensure every derived signal is explainable and references its supporting source type without implying judgement about individuals
- Never store raw audio beyond governed evidence clips (see §5)

4.3 Context Linking (Authoritative)

Findings and drafts MUST be linked to the correct operational context:

- Zone / surgery
- Patient
- Scheduled practitioner
- Scheduled assisting nurse (where available)
- AppointmentID and timestamps

Context linkage is used exclusively for:

- Attaching documentation drafts to the correct appointment
- Routing remediation tasks to the correct role or team
- Preserving traceability for governance and clinical record completeness

The module MUST NOT use context linkage to generate individual performance scoring, ranking, league tables, or automated HR actions.

4.4 Attribution (Deterministic Only)

Where named attribution is required, it MUST be deterministic:

- Primary sources (priority order): device / operator login, workstation acknowledgement, badge / access control events
- Fallback (explicitly enabled only): rota duty assignment, recorded as `rota-attributed`
- Low-confidence cases MUST be recorded as `Unattributed` and MAY raise an admin confirmation task

Named attribution views are restricted to authorised governance roles only; team-level views are the default for all other roles.

All named attribution access is fully audited.

5. Evidence Model (Authoritative)

5.1 Evidence Retention Conditions

Evidence is retained only when both conditions are met:

1. A specific quality finding has been generated
2. The finding requires human validation

No finding → no retained evidence.

5.2 Evidence Scope & Clipping

When retention is triggered:

- Evidence is limited to short, time-bounded excerpts
- Clips are typically bounded to ± 60 –90 seconds around the detected event.
- Full appointments, shifts, or sessions are never retained as evidence

5.3 Evidence Access (Role-Gated, Authoritative)

Evidence:

- Is read-only
- Cannot be downloaded or exported
- Cannot be reused outside the originating quality finding

Access is restricted to:

- Practice owner
- Designated clinical governance lead
- Authorised compliance roles

Evidence is never visible on general staff dashboards.

5.4 Retention & Expiry

- Evidence expires automatically after a defined retention window (e.g., 7–14 days — see §15 Open Questions)
- Extended retention requires a logged governance decision
- Expiry is enforced by the platform; expired evidence cannot be recovered

5.5 Raw Audio Storage Rule (Authoritative)

Default posture is privacy-first: no raw audio is stored. Short evidence clips may be retained only under the conditions in §5.1–5.4. This rule is non-negotiable.

6. Zone-Specific Capabilities

6.1 Clinical Zones (Surgeries)

The module MAY:

- Align capture to the appointment record
- Benchmark against best-practice guides and protocol checklists (where configured)
- Generate draft clinical summaries for clinician review (review-first)
- Flag documentation gaps and critical protocol misses
- Surface comfort / sentiment flags post-appointment only (prompt, not judgement)
- Generate decision-support drafts: treatment plan structure, care plan suggestions, hygiene plan suggestions (all review-first)
- Surface patient-specific details and decisions that should be recorded (review-first)
- Generate patient-safe summaries for the patient app where enabled by practice policy (review-first, clinician-approved)

- Capture and surface patient decisions that materially diverge from practitioner advice as factual record prompts (review-first)
- Create remediation tasks for follow-through

The module **MUST NOT**:

- Auto-finalise any clinical summary or decision-support draft

6.2 Reception / Front-of-House

The module **MAY**:

- Confirm policy recital consistency (deposits, cancellation terms)
- Evaluate enquiry follow-through
- Flag missed handovers
- Link outcomes to Communication Hub and Task Manager for closure

6.3 Decontamination Zones

The module **MAY**:

- Track autoclave utilisation, failures, and throughput
- Detect cycle exceptions (short / long cycles, repeat failures)
- Produce daily summaries and exception alerts

Named attribution in this zone is deterministic and RBAC-controlled; low-confidence cases are recorded as `Unattributed`.

6.4 Arrival / Entry Zones

The module **MAY**:

- Track arrival / departure trends using access logs and check-in data
- Flag overwork / underwork variance against rota
- Surface workload and wellbeing indicators at team level

Named individual views in this zone are restricted to authorised managers and are fully audited.

6.5 Back-Office / Admin Zones

The module **MAY** support admin workflow quality assurance and call-quality assurance in explicitly permitted environments.

6.6 Meeting / Conference Room Zones

Where AI Meeting Notes is active in a meeting or conference room, that room **MUST** be configured as an explicitly enabled zone before any monitoring signals are ingested. This zone type governs the quality-assurance signals derived from AI Meeting Notes recordings of internal staff meetings and governance sessions.

The module **MAY**, where this zone type is explicitly enabled:

- Ingest transcription-completion events from AI Meeting Notes as system-event signals and evaluate documentation completeness (e.g., whether a meeting record and action-item list have been produced for a

scheduled governance meeting)

- Surface speaker attribution confidence scores as a derived signal and raise a finding where confidence falls below a configured threshold, prompting human review and correction of the meeting record before it is finalised
- Flag governance-adherence gaps — for example, where a scheduled compliance or clinical governance meeting has no corresponding AI Meeting Notes transcription or approved meeting record linked to it
- Flag action-item extraction anomalies where the derived signal indicates expected follow-through items were not captured or routed to Task Manager
- Generate a quality finding linked to the meeting context ID (not to an AppointmentID) where any of the above signals indicate an exception
- Create remediation tasks routed via Task Manager for human resolution of identified meeting-record gaps

The module MUST NOT:

- Use meeting zone signals to assess, score, or rank individual contributors to a meeting
- Retain raw audio from meeting zones beyond governed evidence clips under §5
- Surface named speaker attribution data outside authorised governance roles, consistent with §4.4 and §9
- Auto-finalise or auto-commit any meeting record, action-item list, or governance artefact produced in conjunction with AI Meeting Notes

Attribution for meeting zone findings follows the deterministic rules in §4.4; where speaker confidence is insufficient to make a deterministic attribution, the finding MUST be recorded as `Unattributed`. Named attribution views for meeting zones are restricted to authorised governance roles and are fully audited.

7. AI Boundaries (Non-Negotiable)

AI MAY:

- Generate draft clinical summaries, treatment plan structures, care plan suggestions, hygiene plan suggestions, and patient-safe summaries — all review-first
- Highlight patient-specific factors or decisions that should be recorded in the patient record — for clinician review only
- Flag documentation gaps, protocol misses, and timing anomalies for human inspection
- Suggest remediation tasks for human approval and routing
- Summarise zone-level quality indicators for staff and governance dashboards
- Explain the basis of any finding or draft output (explainability obligation)

AI MAY NOT:

- Auto-finalise, auto-commit, or auto-send any clinical note, treatment plan, care plan, hygiene plan, or patient-facing summary
- Make autonomous clinical decisions or replace clinical judgement
- Generate individual performance scores, rankings, or league tables
- Bypass governance, RBAC, audit, or evidence retention controls
- Operate in any zone or on any signal type that has not been explicitly enabled
- Create hidden findings, drafts, or capture surfaces that are not inspectable by authorised roles

- Take any action on behalf of the practice or a clinician without explicit human approval

8. Audit & Compliance

The system MUST log the following events as immutable, exportable audit records:

- All Quality Finding state transitions, with actor identity and timestamp
- All Draft Output Artefact state transitions, including disposition (`Accepted`, `Edited`, `Rejected`) and any content diff on edit
- All AI suggestions (findings, drafts) generated, including which were accepted, edited, or rejected by humans
- All evidence clip creation, access, and expiry events, with actor and purpose
- All zone configuration changes (enablement, disablement, signal-type changes)
- All named attribution access events
- All governance decisions to extend evidence retention beyond the default window
- All `Dismissed` finding closures, including the mandatory reason field
- All cross-module events emitted to or consumed from Task Manager, Communication Hub, and AI Guardian
- All AI Concierge attribution events recorded as AI Quality Monitor signals
- All AI Meeting Notes system-event signals ingested from meeting / conference room zones, including speaker attribution confidence scores and transcription-completion events

Audit logs MUST be immutable, tamper-evident, and exportable for CQC, DSPT, and UK GDPR inspection purposes.

The module is aligned to DSPT, CQC Regulation 17, and UK GDPR expectations.

9. Access Control

Access is governed by Access Manager roles. The following controls apply:

Action	Role(s)
Configure zones and signal types	Practice Owner, designated Compliance role
View team-level quality dashboards and zone summaries	Authorised clinical and operational staff
View named-attribution findings	Practice Owner, Clinical Governance Lead, authorised Compliance roles
Access evidence clips	Practice Owner, Clinical Governance Lead, authorised Compliance roles
Review and action clinician-facing draft outputs	Assigned clinician (per appointment context)

Approve patient-safe summaries for patient app release	Assigned clinician
Dismiss a finding	Authorised governance role (reason required)
Extend evidence retention	Practice Owner, Clinical Governance Lead (logged governance decision)
Read-only audit log access	Authorised Compliance roles
View named speaker attribution findings from meeting zones	Practice Owner, Clinical Governance Lead, authorised Compliance roles

MFA is required for any action that accesses evidence clips or extends evidence retention, given the sensitivity of those operations.

Named attribution views are fully audited. General staff dashboards never surface evidence or named individual data.

10. Integration Summary

- **Task Manager** — outbound: remediation tasks and follow-up work items created by AI Quality Monitor on finding generation (async event)
- **Communication Hub** — outbound: alerts, governance review threads, and escalation notifications on finding state changes (event)
- **AI Guardian** — outbound: quality indicators and zone-summary telemetry for higher-level operational awareness (async)
- **Access Manager** — inbound: RBAC enforcement for all read, write, configure, and approve actions
- **Appointment Manager** — inbound: appointment context (AppointmentID, patient, practitioner, timestamps) for finding and draft linkage (sync lookup)
- **AI Concierge** — inbound: automated-action events (forms sent, tasks created, voice/call recovery outcomes) ingested as system-event signals and logged as attribution events
- **AI Meeting Notes** — inbound: transcription-completion events, speaker attribution confidence scores, and action-item extraction outcomes from meeting / conference room zones, ingested as system-event signals

11. Integration Contracts

11.1 Inbound (this module consumes from)

From module	What	Contract
Appointment Manager	Appointment context (ID, patient, practitioner, zone, timestamps)	Sync lookup
Access Manager	Role and permission data for RBAC enforcement	Sync

AI Concierge	Automated-action events (forms, tasks, call outcomes)	Async event
AI Meeting Notes	Transcription-completion events, speaker attribution confidence scores, action-item extraction outcomes	Async event
PMS (Dentally and others)	Appointment and clinical record cross-verification	Sync / webhook
Digital Forms	Form completion status for cross-verification	Async event
Aftercare	Aftercare delivery status for cross-verification	Async event
Care Plans	Care plan presence and status for cross-verification	Async event

11.2 Outbound (this module emits to)

To module	What	Contract
Task Manager	Remediation tasks and follow-up work items	Async event
Communication Hub	Alerts, escalation threads, governance notifications	Async event
AI Guardian	Quality indicators, zone summaries, telemetry	Async

11.3 PMS Boundary

The PMS (e.g., Dentally) owns the patient record and is the system of record for finalised clinical notes. AI Quality Monitor cross-verifies system events against the PMS for documentation completeness signals, but **MUST NOT** write to the patient record directly. Finalised content flows to the PMS only after a clinician has reviewed, approved, and committed a draft output through the appropriate clinical workflow surface. AI Quality Monitor owns the draft; the PMS owns the record.

12. Delivery Surfaces & Access (Authoritative)

12.1 Web Portal (Staff)

AI Quality Monitor surfaces in the staff web portal as:

- Quality dashboard: zone summaries, finding lists, telemetry indicators (role-gated)

- Finding detail views with explainability notes and (where applicable) evidence access (governance roles only)
- Draft output review queue for clinicians (clinical summary, treatment plan, care plan, hygiene plan drafts)
- Zone and signal configuration screens (admin / compliance roles only)

12.2 Tablet App

Clinician-facing draft outputs (clinical summary, decision-support drafts) are surfaced in the clinician's appointment / day-list view on the tablet, enabling review and approval at point of care.

12.3 Patient Mobile App

Where enabled by practice policy and approved by the clinician, the patient-safe appointment summary is surfaced in the patient app. Content is released only after explicit clinician approval (see §7 AI Boundaries).

12.4 Engagement Signals

AI Quality Monitor emits:

- Zone-level quality indicators to Smart Dashboards (where available)
- Finding and remediation telemetry to AI Guardian
- Remediation task completion signals back from Task Manager (consumed to update finding state)

13. Non-Functional Requirements

- **Performance:** Finding generation and draft output creation **MUST** complete within a latency envelope that does not disrupt clinical workflow; specific SLA targets to be defined during engineering design (see §15).
- **Reliability:** The module **MUST** degrade gracefully if upstream signal feeds are unavailable — no finding should be generated from incomplete or corrupted signals; the absence of a signal must be distinguishable from a detected gap.
- **Scalability:** The module **MUST** support multi-site, multi-zone, multi-tenant deployments with per-zone configuration isolation.
- **Security:** All patient-bound data and evidence clips **MUST** be encrypted at rest and in transit. Key management and secrets handling follow platform-level security policy (Platform Security & Governance). Evidence clips are stored in access-controlled, isolated storage with no public endpoints.
- **Privacy:** The module honours UK GDPR rights including subject access, erasure (within the constraints of clinical retention obligations), and purpose limitation. Raw audio is not retained beyond governed evidence clips. Data retention windows for findings and evidence are enforced by the platform.
- **Accessibility:** All staff-facing review surfaces **MUST** meet WCAG 2.1 AA accessibility standards.
- **Observability:** The module **MUST** export: finding generation rate by zone, draft output acceptance / rejection rates, evidence clip creation and expiry counts, remediation task creation and completion rates, signal ingestion error rates. Traces **MUST** be available for finding-generation pipelines to support debugging without exposing patient data.

14. Build Contract (Engineering & QA)

14.1 Canonical Data Model

(Canonical schema field names and table structures to be defined during engineering design. The following objects are normatively established in §3 and must be implemented: `quality_finding`, `draft_output_artefact`, `evidence_clip`, `zone_configuration`, `derived_signal`, `attribution_event`.)

14.2 Core Behaviour Rules

The following rules are testable and must be implemented by engineering and verified by QA:

1. No finding is generated without an explainability note referencing the derived signal type that triggered it.
2. No evidence clip is created unless a quality finding exists and requires validation (§5.1).
3. No evidence clip is accessible to any user not in an authorised governance role.
4. No Draft Output Artefact transitions to `Accepted` or `Edited` without an explicit human action by an authorised user — no automated commit path exists.
5. No raw audio is stored beyond governed evidence clips.
6. All state transitions on Quality Findings and Draft Output Artefacts are logged with actor identity and timestamp.
7. All named attribution views are restricted to authorised governance roles; team-level views are the default.
8. A finding in `Dismissed` state requires a non-null reason field; the system MUST reject dismissal submissions without one.
9. Evidence clips expire automatically at the end of the defined retention window; expiry cannot be reversed except via a logged governance decision.
10. Zone monitoring MUST NOT activate on any signal type not explicitly enabled in zone configuration.
11. AI Concierge events are recorded as system-event signals and logged as attribution events regardless of whether they originate in a physical zone.
12. Patient-safe summaries MUST NOT be released to the patient app without explicit clinician approval (disposition `Accepted` or `Edited` on the relevant Draft Output Artefact).
13. AI Meeting Notes events (transcription-completion, speaker attribution confidence, action-item extraction) MUST only be ingested as quality signals where the meeting / conference room zone is explicitly enabled; findings generated from these signals MUST carry a meeting context ID rather than an AppointmentID.
14. Speaker attribution confidence scores from AI Meeting Notes MUST be recorded as `Unattributed` on any resulting finding where confidence falls below the configured threshold; named speaker attribution views for meeting zone findings are restricted to authorised governance roles.

14.3 Configuration Surfaces

- **Practice-level (Admin Control Plane):** Zone enablement, signal type activation per zone, evidence retention window (within platform bounds), named attribution enablement per zone, practice policy flags (e.g., patient-safe summary enabled, meeting / conference room zone enabled, speaker attribution confidence threshold)
- **Role-level (Access Manager):** Which roles can access named attribution views, evidence clips, and zone configuration screens
- **Per-finding overrides (this module):** Extended retention decision (logged governance action), dismissal reason

14.4 Filtering & Views

Standard filters the UI must support:

- By zone
- By finding type
- By finding state
- By date range
- By attribution status (`Named` / `Rota-attributed` / `Unattributed`)
- By draft output type and disposition
- By appointment / context ID

Saved views must support at minimum: open findings by zone, pending draft reviews by clinician, remediation task backlog.

14.5 Module Extension Map

Future extensions must not break the evidence model (§5), state machines (§3.2, §3.4), or AI boundaries (§7). Extension points include:

- Additional zone types (additive — new zone configuration records only)
- Additional draft output types (additive — new `DraftType` enum values and review surfaces)
- Additional derived signal types (additive — new signal definitions; existing finding logic is unaffected)
- Additional telemetry exports (additive — new metric definitions; existing observability contracts unchanged)

14.6 Acceptance Criteria

The build of AI Quality Monitor is complete when:

- Zone-based alignment is accurate: findings and drafts are correctly linked to appointment ID, call thread, sterilisation cycle ID, meeting context ID, or rota window as appropriate
- All Quality Finding and Draft Output Artefact state transitions enforce rules in §3.2 and §3.4
- Findings are explainable: every finding carries an explainability note referencing the derived signal type
- Evidence clips are governed: clipped, time-limited, RBAC-controlled, and expiry-enforced per §5
- All Draft Output Artefact types in §7 are implemented as review-first workflows; no auto-commit path exists
- Decision-support drafts (treatment plan, care plan, hygiene plan) are review-first and never auto-finalised
- Remediation tasks and alerts route correctly via Task Manager and Communication Hub
- Privacy-first posture holds: no raw audio is stored beyond governed evidence clips
- Monitoring cannot be used as punitive surveillance by default: no individual scoring or ranking surface exists
- AI boundaries in §7 are enforced; negative tests (auto-commit, unapproved evidence access, scoring) pass
- Audit log captures every event in §8; logs are immutable and exportable
- Access control is enforced per §9; named attribution restricted to governance roles
- All non-functional requirements in §13 are met
- AI Concierge events are correctly ingested, attributed, and logged
- AI Meeting Notes events are correctly ingested from explicitly enabled meeting / conference room zones; speaker attribution confidence findings are raised and restricted per §6.6 and §4.4; no meeting zone

monitoring activates without explicit zone enablement

15. Versioning & Governance

This specification is owned by: the AI Quality suite module owner.

Changes to this spec require:

- Review by the Post-MVP module owner
- Impact analysis across declared related modules (Task Manager, Communication Hub, AI Guardian, Access Manager, Appointment Manager, AI Concierge, AI Meeting Notes)
- Version bump (patch for clarifications; minor for capability additions; major for state machine, evidence model, or AI boundary changes)

16. Explicit Non-Goals

- Individual staff performance scoring or ranking — prohibited; no module currently owns this and it must not be added to AI Quality Monitor
- Continuous video recording or surveillance — prohibited; camera inputs are event-feeds only
- Autonomous clinical decisions or auto-finalised clinical records — prohibited; owned by the clinician via the PMS
- Disciplinary evidence generation by default — out of scope; any governance extension requires a separate, logged governance decision and is not part of this module's core capability

17. Open Questions

These questions are present-but-unresolved in the source material and must be decided before this spec can be promoted from `draft` to `published`.

1. **Evidence retention window:** The original states "e.g., 7–14 days" as an example range. What is the authoritative default retention window for evidence clips, and what is the maximum permitted extension window for governance-approved extensions?
2. **Performance SLAs:** No latency or throughput targets are defined for finding generation or draft output creation. What are the acceptable latency envelopes, particularly for clinical-zone outputs that must not disrupt appointment flow?
3. **AI Concierge boundary:** The spec states AI Concierge events are logged as AI Quality Monitor attribution events. Is AI Concierge responsible for emitting a well-defined event contract to AI Quality Monitor, or does AI Quality Monitor poll / consume a shared event stream? The integration contract type needs to be confirmed.
4. **Patient-safe summary enablement:** The spec states this feature is enabled "where enabled by practice policy." Does this require explicit opt-in at practice level, and is there a platform-level default (enabled or disabled)?
5. **Sentiment flags scope:** Clinical zones may surface "comfort/sentiment flags post-appointment only." The basis and permissible signal types for sentiment detection are not defined. This needs scoping before build to ensure it does not inadvertently cross the individual-scoring prohibition.
6. **AI Meeting Notes event contract:** AI Meeting Notes emits transcription-completion, speaker attribution confidence, and action-item extraction events to AI Quality Monitor. The precise event schema, confidence-score thresholds that trigger a finding, and whether AI Meeting Notes or AI Quality Monitor is

responsible for defining those thresholds need to be agreed before build of the meeting / conference room zone capability.