

Aftercare Manager

Doc type: technical · **Version:** v0.1 · **Status:** published · **Module slug:** aftercare-manager
Exported: 2026-05-15 11:12 UTC · **By:** anonymous

Aftercare Manager – Technical Specification

1. Module Purpose & Scope (Authoritative)

Aftercare Manager is the single system of record for post-treatment guidance, follow-ups, and recovery support within Primoro. It closes the care loop after an appointment ends, turning aftercare into a governed, proactive stage of the patient journey rather than an afterthought. The module ensures that patients receive clear, treatment-linked instructions without relying on memory, paper, or unmanaged PDFs; that potential issues are identified earlier and not reactively; and that clinicians and practices are protected by audit-ready records.

It governs:

- The full lifecycle of Aftercare Instructions, from template selection through delivery, patient interaction, escalation, and resolution.
- The Aftercare Template Library, including versioning, treatment-type mappings, and delivery rules.
- Automated and AI-assisted follow-up, check-in, and escalation workflows.

It explicitly does not:

- Own appointment booking or diary control — governed by Appointment Manager.
- Own free-form staff messaging or general patient chat — governed by Communication Hub.
- Own task execution — governed by Task Manager.
- Author or store clinical records — owned by the PMS.
- Manage payments, membership plans, or reputation/feedback — governed by their respective modules.

2. Ownership & Responsibilities

2.1 Aftercare Manager IS Responsible For

- Linking aftercare instructions to completed treatments and maintaining those links immutably.
- Managing a versioned aftercare template library, including out-of-the-box templates and practice-level customisation.
- Delivering aftercare digitally via the patient mobile app and maintaining persistent patient access to delivered instructions.
- Running automated follow-up check-ins and routing patient responses through AI-first interaction (Aiden) with human fallback.
- Escalating concerns to the practice team with full conversation context preserved.
- Supporting delivery to family members or carers via Family Profiles, with dual attribution on any escalated record (PatientId and CarerId as distinct fields).

- Maintaining immutable audit trails covering delivery, access, patient responses, AI participation, escalation events, and staff responses.
- Emitting higher-risk patient pattern signals to AI Guardian (when enabled). For the purposes of this boundary, a **higher-risk pattern signal** is any of the following: an unanswered follow-up that has exceeded the configured overdue threshold; a patient response that has triggered a negative or concerning sentiment flag; a keyword or symptom pattern match against the configured escalation keyword set; or a repeated escalation for the same patient or treatment type within a configured time window. These signals are emitted as structured events and **MUST NOT** include unstructured clinical narrative. Ownership of any Finding raised from these signals resides with AI Guardian; Aftercare Manager retains ownership only of the underlying aftercare state data.

2.2 Aftercare Manager IS NOT Responsible For

- Booking or modifying appointments triggered by escalation — owned by Appointment Manager.
- Executing or tracking staff follow-up tasks — owned by Task Manager.
- Sending general practice communications outside the aftercare thread — owned by Communication Hub.
- Storing or modifying clinical notes — owned by the PMS.
- Performing risk-pattern analysis across the patient population — owned by AI Guardian (when enabled).
- Enforcing identity, authentication, or role definitions — owned by Access Manager.

3. Core Objects (Normative)

3.1 Aftercare Instruction (Canonical Artefact)

An Aftercare Instruction is a governed digital artefact representing structured post-treatment guidance delivered to a patient (or their carer) following a completed treatment.

It is treatment-linked, versioned, delivered digitally, readable at any time by the patient, and fully auditable.

Minimum required fields:

- AftercareId (immutable)
- TemplateId
- TemplateVersion
- TreatmentTypeId
- Origin (Appointment | Manual | Bulk)
- PatientId
- AppointmentId (nullable for Manual and Bulk origins)
- CarerId (nullable)
- AftercareState
- DeliveredTimestamp
- ViewedTimestamp (nullable until viewed)
- FollowUpSchedule (nullable)
- ConversationId (link to Communication Hub thread)
- Escalated (boolean)

- EscalationTimestamp (nullable)
- CreatedBy (user/role or system)
- CreatedAt
- AuditLog (immutable append-only)

3.2 Aftercare Instruction State Machine (Authoritative)

States:

- **Prepared** — template selected; instruction not yet delivered.
- **Delivered** — instruction sent to patient or carer; delivery logged.
- **Viewed** — patient (or carer) has opened the instruction; timestamp recorded.
- **Follow-Up Triggered** — automated check-in dispatched per configured schedule.
- **AI-Handled** — patient interaction in progress via Aiden; no human intervention yet required.
- **Escalated** — conversation routed to practice team; full context preserved.
- **Resolved** — concern addressed; record closed.

Rules:

- All state transitions are auditable and time-stamped.
- An Aftercare Instruction **MUST NOT** return to Prepared once it reaches Delivered.
- Escalated **MUST NOT** be set without preserving the full prior conversation context.
- Only authorised staff roles (as defined in Access Manager) may manually trigger transitions from Prepared to Delivered, or mark a record Resolved.
- A patient or carer interaction alone **MUST NOT** close (Resolve) an escalated record; a named staff member must act.

3.3 Aftercare Template (Canonical Artefact)

An Aftercare Template defines reusable aftercare content, the treatment types it applies to, delivery rules, and follow-up behaviour.

Minimum required fields:

- TemplateId (immutable)
- TemplateVersion
- TreatmentTypeIds (one or many)
- ContentBody (structured sections)
- DeliveryRules
- FollowUpBehaviour (nullable)
- CreatedBy
- CreatedAt
- AuditLog (immutable)

Templates are version-controlled. Prior versions **MUST** be retained for audit purposes and **MUST NOT** be overwritten destructively.

4. Creation, Trigger Sources & Delivery

4.1 Appointment-Driven Trigger (Primary) (Authoritative)

Aftercare is triggered automatically when an appointment is marked Completed in Appointment Manager.

The module **MUST**:

- Identify the relevant treatment type(s) from the completed appointment.
- Select the matching aftercare template.
- Deliver the Aftercare Instruction to the patient immediately.
- Log delivery with a timestamp regardless of whether the patient opens the instruction.

The module **MUST NOT**:

- Send aftercare if no applicable template exists. In this case, the absence of a matched template **MUST** be logged as a non-delivery event against the appointment record so that staff are aware.

4.2 Manual Issuance

Authorised staff **MAY**:

- Issue aftercare instructions manually outside the appointment-completion flow.
- Re-issue updated instructions when clinically justified, with the re-issuance logged as a distinct event.
- Attach supplementary guidance to an existing Aftercare Instruction.

All manual issuance **MUST** be attributed to the acting staff member and logged.

4.3 Bulk Issuance

Aftercare **MAY** be issued in bulk to defined patient cohorts, for audit or safety reviews, or for group-wide standards in multi-site operations.

All bulk issuance actions **MUST** be auditable, including the cohort definition, the template version used, and the authorising staff member.

4.4 Task Manager Boundary for Aftercare-Triggered Staff Tasks

Aftercare Manager does not create general-purpose tasks in Task Manager as part of the instruction delivery flow. Delivery of an Aftercare Instruction — whether appointment-driven, manual, or bulk — is handled entirely within this module and does not generate a corresponding Task Manager task.

Task Manager **IS** invoked by Aftercare Manager in two specific, bounded circumstances:

- **On escalation** — when an Aftercare Instruction transitions to the Escalated state, Aftercare Manager emits an outbound event to Task Manager to generate a staff follow-up task. This task carries the AftercareId and a reference to the escalated conversation, so the responding staff member is directed to the correct context. The task is owned and tracked by Task Manager; Aftercare Manager does not duplicate task-tracking internally.
- **When AI Guardian raises a Finding referencing an AftercareId** — Task Manager **MAY** receive a task from AI Guardian for staff review. Aftercare Manager plays no role in that secondary task-creation flow; it has already fulfilled its obligation by emitting the underlying signal.

All other staff actions related to aftercare (template management, manual issuance, resolution of escalations) are performed directly in Aftercare Manager's staff portal and do not route through Task Manager.

4.5 AI-Assisted Patient Support (Where Enabled) (Authoritative)

Aiden (Primoro AI Assistant) MAY:

- Answer common aftercare questions using the exact instructions delivered to that patient.
- Provide reassurance and structured guidance drawn from the delivered aftercare and approved recovery knowledge.
- Monitor patient responses for sentiment, keywords, and symptom patterns and flag conditions that meet escalation thresholds.

Aiden MUST NOT:

- Provide clinical diagnoses or prescriptions.
- Override clinician instructions embedded in the delivered aftercare.
- Make autonomous decisions without triggering escalation when thresholds are met.
- Close or resolve a conversation silently without logging the outcome.

5. Delivery Surfaces & Access (Authoritative)

5.1 Patient Mobile App

- Aftercare Instructions appear as a persistent, readable record accessible at any time after delivery.
- Follow-up check-ins and AI-assisted interactions surface within the same conversational thread as the original instruction.
- Patients can ask questions and raise concerns without leaving the aftercare context.
- Where the patient is a dependant and delivery is via a carer, the carer's Family Profiles session governs visibility; dual attribution (PatientId and CarerId) is applied to all interactions.

5.2 Guardian (Carer) Delegated Access

Where a patient is a dependant managed through Family Profiles, a guardian MAY access that dependant's Aftercare Instructions within the patient mobile app under the following rules:

- The guardian's access is scoped strictly to dependants linked to their account in Family Profiles. A guardian MUST NOT have visibility of Aftercare Instructions for patients outside their authorised dependant set.
- The guardian experiences the same persistent-record access as a patient: delivered instructions are readable at any time, and follow-up check-ins and AI-assisted interactions surface in the same conversational thread.
- AI-first interaction (Aiden) MAY respond to guardian queries about a dependant's aftercare using the same content boundary that applies to direct patient interactions — the specific delivered instruction and approved recovery knowledge only.
- Escalation triggered during a guardian-led interaction follows the same automated threshold logic as patient-direct escalation. The escalated record MUST carry both PatientId and CarerId as distinct, non-collapsed fields so that the responding staff member can identify both the patient and the acting guardian unambiguously.

- Guardian identity MUST NOT overwrite or be merged into the patient record at any point in the delivery, interaction, or audit trail. The distinction between the patient as the subject of care and the guardian as the acting user MUST be preserved throughout.
- If a guardian's delegated authority is revoked in Family Profiles, access to the dependant's Aftercare Instructions MUST be withdrawn in real time. Aftercare Manager MUST honour this revocation without requiring a separate configuration action.

5.3 Staff Web Portal

- Authorised staff can manage the aftercare template library, including creating, editing (with version bump), and mapping templates to treatment types.
- Live dashboards display follow-up status, escalation state, and AI vs. human involvement for active Aftercare Instructions.
- Escalated conversations surface with full context — prior AI interaction, patient responses, and treatment history — so staff can respond as named individuals.
- AI Guardian findings that reference an AftercareId (when AI Guardian is enabled) are surfaced alongside the relevant instruction record.

5.4 Tablet App

- Read-only access to active Aftercare Instructions where tablet access is permitted by role.
- Supports in-surgery reference to aftercare guidance if clinically required.

5.5 Engagement Signals

The module emits the following signals for staff visibility and analytics:

- Delivery confirmation per instruction.
- Patient view events.
- Follow-up response rates.
- Escalation frequency by treatment type, staff routing, and time-to-resolution.
- AI-handling rate vs. escalation rate, surfaced as an operational metric.

6. Integration Contracts

6.1 Inbound (this module consumes from)

From module	What	Contract
Appointment Manager	Appointment Completed event with treatment type(s)	Async event
Access Manager	Role and permission grants for staff and patient access	Sync query at access time
Family Profiles	Carer-patient relationships and active guardian sessions	Sync query at delivery time

Communication Hub	Conversation thread state and message delivery receipts	Async event
-------------------	---	-------------

6.2 Outbound (this module emits to)

To module	What	Contract
Communication Hub	Aftercare delivery, follow-up messages, escalation thread creation	Event
Task Manager	Staff follow-up tasks on escalation	Event
Appointment Manager	Follow-up appointment recommendation on escalation	Event
AI Guardian	Aftercare state signals (unanswered follow-ups, sentiment flags, keyword matches, repeated escalations)	Event (when AI Guardian is enabled)
Audit & Compliance	All state-transition and interaction audit events	Immutable event log

6.3 PMS Boundary

The PMS is the system of record for clinical notes, diagnoses, and treatment authoring. Aftercare Manager consumes treatment type identifiers from the PMS (via Appointment Manager) as triggers for template selection, but **MUST NOT** write to or modify any clinical record held in the PMS. Aftercare instructions are a governed Primoro artefact, not a PMS document.

7. AI Boundaries (Non-Negotiable)

Module does embed an AI surface: Aiden (Primoro AI Assistant) operates as first-line responder to patient aftercare questions.

AI MAY:

- Answer patient questions using only the aftercare instructions delivered to that specific patient and approved recovery knowledge.
- Provide reassurance and structured guidance within the scope of delivered content.
- Monitor sentiment and keywords to identify escalation conditions.
- Summarise AI-handled interactions for staff review.

AI MAY NOT:

- Give clinical diagnoses or prescriptions.
- Override or contradict instructions authored by the clinician.

- Make autonomous decisions on behalf of the practice; all decisions that exceed confidence thresholds MUST trigger escalation to a human.
- Close or resolve an escalated conversation without documented human action.
- Bypass governance, audit, or access controls.
- Replace required clinical judgement.

8. Audit & Compliance

The system MUST log:

- Aftercare Instruction creation, with origin (Appointment | Manual | Bulk), template version, and acting user or system.
- Delivery events, including timestamp and recipient identity (PatientId and CarerId where applicable).
- Patient and carer view events with timestamp.
- Follow-up dispatch and patient response events.
- All Aiden (AI) interactions: inputs, outputs, and whether the response was accepted, escalated, or overridden by a human.
- Escalation events: trigger reason (patient request | confidence threshold | keyword/sentiment flag), timestamp, and staff member to whom the conversation was routed.
- Staff responses within escalated conversations, attributed to the named individual.
- Template create, edit, and version-bump events.
- Bulk issuance events, including cohort definition and authorising staff member.
- Guardian-delegated interactions, with dual attribution (PatientId and CarerId) preserved as distinct fields.
- Non-delivery events where no template matched a completed appointment.
- Guardian access-revocation events received from Family Profiles, with timestamp and the resulting withdrawal of access to the dependant's records.

All audit logs MUST be immutable (append-only) and exportable for regulatory inspection. Guardian identity MUST NOT overwrite or be collapsed into the patient record at any point in the audit trail.

9. Access Control

Access is governed by Access Manager using role-based permissions.

Action	Who may perform it
View own Aftercare Instructions	Patient (own records only)
View Aftercare Instructions as carer	Family Profiles-authorized guardian (dependant's records only)
View all instructions (staff)	Authorised clinical and admin roles per Access Manager
Create / issue Aftercare Instructions	Authorised staff roles

Manage Aftercare Templates	Admin roles
Respond to escalated conversations	Authorised clinical roles
Resolve escalated instructions	Authorised clinical roles
View audit logs	Compliance / admin roles

- Patient access is strictly limited to their own records.
- Carer access is limited to dependants linked via Family Profiles.
- Staff access is revoked immediately on exit from the practice, enforced by Access Manager.
- MFA requirements for access to escalated patient communications and audit log export are governed by Access Manager policy and MUST be honoured by this module.

10. Integration Summary

- **Appointment Manager** — inbound Appointment Completed events that trigger automatic aftercare delivery; outbound follow-up appointment recommendations on escalation.
- **Communication Hub** — delivery of aftercare messages, follow-up threads, and escalation conversations.
- **Task Manager** — outbound staff follow-up tasks generated on escalation only; aftercare delivery itself does not generate tasks.
- **Family Profiles** — carer-patient relationship data for dependant delivery, delegated guardian access, and dual attribution; guardian access revocations are consumed in real time.
- **Access Manager** — RBAC for all read, write, and approve operations; MFA policy enforcement.
- **AI Guardian** — outbound aftercare state signals (unanswered follow-ups, sentiment flags, keyword matches, repeated escalations) for population-level risk detection (when enabled); Finding ownership resides with AI Guardian.
- **Audit & Compliance** — immutable event log for all state transitions and interactions.
- **PMS** — treatment type identifiers consumed as triggers; no write-back to PMS.

11. Explicit Non-Goals

- Clinical diagnosis or prescription — MUST NOT be performed by this module or by Aiden acting within it; these remain within the PMS and clinician scope.
- Reputation management or patient feedback collection — if added to the platform, this would be owned by a dedicated Feedback module.
- Autonomous AI decision-making without human oversight — prohibited; Aiden remains assistive only.
- General staff-to-patient messaging outside the aftercare context — owned by Communication Hub.
- Unmanaged PDF or portal-based delivery — explicitly out of scope; delivery is always through the patient mobile app.

12. Versioning & Governance

This specification is owned by: the Aftercare Manager module owner.

Changes to this spec require:

- Review by the MVP module owner.
- Impact analysis across all declared related modules (see /propose).
- Version bump: patch for clarifications, minor for new capabilities, major for breaking contract changes.
- All changes MUST preserve auditability, keep AI assistive rather than authoritative, and maintain separation of concerns aligned with Primoro CORE principles.

13. Build Contract (Engineering & QA)

13.1 Canonical Data Model

Aftercare Instruction table (minimum schema):

```
aftercare_instruction (
  aftercare_id          UUID PRIMARY KEY,
  template_id          UUID NOT NULL,
  template_version     TEXT NOT NULL,
  treatment_type_id    UUID NOT NULL,
  origin               ENUM('Appointment', 'Manual', 'Bulk') NOT NULL,
  patient_id          UUID NOT NULL,
  appointment_id       UUID NULLABLE,
  carer_id             UUID NULLABLE,
  aftercare_state      ENUM('Prepared', 'Delivered', 'Viewed',
                            'FollowUpTriggered', 'AIHandled',
                            'Escalated', 'Resolved') NOT NULL,
  delivered_timestamp  TIMESTAMPTZ NULLABLE,
  viewed_timestamp     TIMESTAMPTZ NULLABLE,
  follow_up_schedule   JSONB NULLABLE,
  conversation_id      UUID NULLABLE,
  escalated            BOOLEAN NOT NULL DEFAULT FALSE,
  escalation_timestamp TIMESTAMPTZ NULLABLE,
  created_by          TEXT NOT NULL,
  created_at          TIMESTAMPTZ NOT NULL,
  audit_log           JSONB NOT NULL
)
```

Aftercare Template table (minimum schema):

```
aftercare_template (
  template_id          UUID NOT NULL,
  template_version     TEXT NOT NULL,
  treatment_type_ids   UUID[] NOT NULL,
  content_body         JSONB NOT NULL,
  delivery_rules       JSONB NOT NULL,
  follow_up_behaviour  JSONB NULLABLE,
  created_by          TEXT NOT NULL,
  created_at          TIMESTAMPTZ NOT NULL,
  audit_log           JSONB NOT NULL,
  PRIMARY KEY (template_id, template_version)
)
```

13.2 Core Behaviour Rules

1. An Aftercare Instruction **MUST** be created and delivered automatically when an Appointment Completed event is received and a matching template exists for the treatment type.
2. If no matching template exists, a non-delivery event **MUST** be logged against the appointment; no silent failure is permitted.
3. Patients **MUST** be able to re-read their Aftercare Instruction indefinitely after delivery; access **MUST NOT** be revoked except under documented data-retention policy.
4. All Aiden (AI) responses **MUST** reference only the content of the delivered Aftercare Instruction and approved recovery knowledge; responses outside this scope are prohibited.
5. Escalation **MUST** fire automatically when AI confidence thresholds, keyword matches, or sentiment flags are exceeded; human action is required to resolve the escalated record.
6. Escalated records **MUST** carry the full prior conversation context; context **MUST NOT** be truncated or discarded on escalation.
7. Where a guardian (carer) is the interacting user, escalated records **MUST** carry both PatientId and CarerId as distinct fields; the guardian identity **MUST NOT** overwrite the patient record.
8. AI Guardian signals **MUST** be emitted for: unanswered follow-ups beyond the configurable overdue threshold; negative/concerning sentiment; flagged keywords or symptom patterns; and repeated escalations for the same patient or treatment type within a configured window.
9. All state transitions **MUST** be time-stamped and appended to the immutable AuditLog; no transition may be deleted or modified.
10. Staff access **MUST** be revoked in real time on notification from Access Manager; revoked staff **MUST NOT** be able to read, write, or respond to any Aftercare record.
11. Template edits **MUST** create a new TemplateVersion; prior versions **MUST** be retained and **MUST NOT** be overwritten.
12. Guardian access revocations received from Family Profiles **MUST** withdraw access to the dependant's Aftercare Instructions in real time; no additional configuration action by staff is required to enforce revocation.
13. Task Manager **MUST** be invoked via outbound event only on escalation; aftercare instruction delivery (appointment-driven, manual, or bulk) **MUST NOT** generate a task in Task Manager.

13.3 Configuration Surfaces

Practice-level (Admin Control Plane): - Aftercare template content and versioning. -

Treatment-type-to-template mappings. - Follow-up timing rules per treatment type. - Escalation thresholds (AI confidence, sentiment, keywords). - Staff routing rules for escalated conversations. - Family/carers delivery settings. - Overdue follow-up threshold for AI Guardian signal emission.

Per-user preferences (Access Manager): - Notification preferences for escalation routing (within role-permitted scope).

Per-instruction overrides (this module): - Manual issuance of supplementary guidance by authorised staff.

13.4 Filtering & Views

Staff views **MUST** support filtering by:

- Treatment type.
- Delivery status (Prepared | Delivered | Viewed).
- Follow-up status (scheduled | triggered | responded | overdue).

- Escalation state (not escalated | escalated | resolved).
- AI vs. human involvement in conversation handling.
- Date range of delivery or escalation.
- Origin (Appointment | Manual | Bulk).

13.5 Module Extension Map

The following extension points are defined for future capability without breaking this contract:

- **AI Guardian** (when enabled) — Aftercare Manager exposes aftercare state signals as inputs (unanswered follow-ups, sentiment flags, keyword matches, repeated escalations). AI Guardian MAY raise risk alerts, recovery-concern summaries, and follow-up recommendations from these signals. Every finding that references an AftercareId MUST be routed via Task Manager or Communication Hub and MUST be surfaced in the Aftercare Manager staff view alongside the relevant record.
- **Governance Reporting** (when enabled) — compliance summaries and aftercare delivery metrics.
- **Group Controls** (when enabled) — multi-site template consistency and group-wide issuance.

13.6 Acceptance Criteria

The build of Aftercare Manager is complete when:

- [] Aftercare Instructions are created and delivered automatically on Appointment Completed events where a matching template exists.
- [] Non-delivery is logged (not silently dropped) when no template matches.
- [] Patients can re-read delivered instructions at any time via the patient mobile app.
- [] Template versioning is enforced; prior versions are retained.
- [] Manual and bulk issuance flows are attributable and auditable.
- [] AI (Aiden) operates within approved content boundaries; out-of-scope responses are blocked (negative tests pass).
- [] Escalation triggers automatically on threshold breach; full context is preserved on escalation.
- [] Guardian-delegated interactions carry dual attribution (PatientId + CarerId) throughout, including in audit log.
- [] Guardian access revocations from Family Profiles withdraw access in real time without requiring manual staff intervention (verified by integration test).
- [] All state transitions are logged to the immutable AuditLog with actor and timestamp.
- [] Staff access is revoked in real time per Access Manager instruction.
- [] AI Guardian signal events are emitted correctly for all four defined signal types (when AI Guardian is enabled).
- [] Task Manager receives an outbound escalation event on each escalation transition, and no task is generated for routine delivery events (positive and negative tests pass).
- [] All integrations in §6 are wired and tested.
- [] All non-functional requirements in §14 are met.

14. Non-Functional Requirements

- **Performance:** Aftercare Instruction delivery MUST be initiated within 5 seconds of receipt of an Appointment Completed event. Follow-up dispatch MUST NOT block or delay any other system action. AI-first response to patient messages MUST be acknowledged within a timeframe consistent with conversational UX expectations — target defined at infrastructure layer.
- **Availability:** The delivery and patient-access surfaces MUST be available at a level consistent with patient-facing Primoro platform SLAs. Silent failures in delivery are not acceptable; the module MUST implement retry logic and alert on persistent delivery failure.
- **Reliability:** No Aftercare Instruction may be lost in transit. Delivery failures MUST be detected, logged as non-delivery events, and surfaced to staff. The module MUST degrade gracefully — if AI-first interaction is unavailable, the conversation MUST be routed directly to the human escalation path without patient-visible failure.
- **Scalability:** The module MUST support multi-site, multi-tenant operations. Bulk issuance MUST be designed to avoid degrading system performance for concurrent appointment-triggered deliveries. Template management and cohort queries MUST remain performant as template libraries and patient volumes grow.
- **Security:** All patient data MUST be encrypted in transit (TLS) and at rest. Carer identity data MUST be handled with the same protections as patient data. Access to audit logs MUST be restricted to authorised compliance and admin roles.
- **Privacy:** The module MUST honour patient data-retention policies as configured at the platform level, including the right to erasure where applicable under GDPR, except where retention is required for regulatory or audit purposes. CarerId linkages MUST be subject to the same retention governance as PatientId records.
- **Observability:** The module MUST export: delivery success/failure rates per treatment type; follow-up response rates; escalation rates and time-to-resolution; AI-handling vs. escalation ratio; and audit log write confirmation. All events MUST carry sufficient trace context to support distributed tracing across the Primoro platform.

15. Open Questions

1. **No-template non-delivery notification:** When no template matches a completed appointment, should the non-delivery event generate an automatic staff alert or task, or is passive logging to the audit trail sufficient? *(present in §4.1 as a TBD behaviour — logging is mandated but active notification is unresolved.)*
2. **Follow-up timing rules:** The specification mandates that follow-up timing is configurable but does not define default schedules for out-of-the-box treatment types. What are the default follow-up intervals, and who owns the initial default library? *(present in §13.3 as a configurable surface without defaults.)*
3. **AI confidence thresholds:** Escalation thresholds for AI confidence, sentiment, and keyword detection are described as configurable but no default values or ranges are specified. Who sets initial defaults, and what is the governance process for tuning them? *(present in §6.3 and §13.3 as configurable without definition.)*
4. **Bulk issuance cohort definition:** The specification states bulk issuance targets "defined patient cohorts" but does not define how cohorts are constructed or whether cohort definition is owned by this module or another. *(present in §4.3 without ownership resolution.)*
5. **Template library governance for multi-site operators:** Group Controls is listed as a future extension for multi-site template consistency, but it is unresolved whether individual sites may override group-level templates and under what conditions. *(present in §13.5 without a policy decision.)*