AI Assistant (Aiden)

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

AI Assistant (Aiden) – Technical Specification

1. Module Purpose & Scope (Authoritative)

AI Assistant — Aiden — is Primoro's calm, action-first intelligence layer, embedded across the platform to guide patients and support staff without replacing human judgement or ownership. It helps users understand what they are seeing in plain English, guides them to complete supported actions inside Primoro, and surfaces what is due, missing, or at risk using the CORE calendar. Aiden explains and supports — it does not decide, diagnose, or act autonomously.

It governs:

- Context-aware guidance delivered across all enabled Primoro surfaces (Patient App, Staff App, Web Portal, Website chat, AI Phone Assistant)
- Action-first conversational flows for patients covering appointments, forms, treatment plan explanation (non-clinical), and payments where enabled
- Role-aware operational assistance for staff, including how-to guidance, at-risk work surfacing, and suggested next actions
- Escalation rule enforcement, clinical safety boundaries, and confidence-threshold management
- AI audit logging with mandatory AI involvement flags on every AI-emitted event

It explicitly does not:

- Provide clinical advice, diagnosis, or outcome commentary (no Primoro module owns this; it is out of scope for the platform)
- Own booking availability logic — that is owned by Appointment Manager
- Execute or own tasks — that is owned by Task Manager
- Own message routing or SLA management — that is owned by Communication Hub
- Manage staff availability decisions — that is owned by Rota Manager
- Produce business analytics or reporting — that is owned by Smart Dashboards

---

2. Ownership & Responsibilities

2.1 AI Assistant IS Responsible For

- Providing context-aware, screen-aware guidance across all enabled Primoro surfaces
- Guiding patients through action-first flows: appointments (within rules), forms and acknowledgements, treatment plan explanation (non-clinical), and balances/payments (where enabled)

- Providing staff with role-aware how-to guidance, operational clarity, suggested next actions, and internal Q&A from practice-approved knowledge
- Surfacing and contextualising alert signals from Smart Dashboards without re-evaluating or contradicting those signals
- Highlighting potential unlinked dependants and prompting staff to consider Family Profile creation — Aiden MAY only highlight and prompt; it MUST NOT create family links, assign permissions, or override consent
- Enforcing escalation rules, clinical safety boundaries, and confidence thresholds consistently across all channels
- Respecting Access Manager RBAC and data-governance rules at all times
- Logging all AI activity with an `ai_involved` flag and `DeviceId` for audit and continuous improvement

2.2 AI Assistant IS NOT Responsible For

- Clinical advice, diagnosis, or recommendations — out of scope for the platform
- Booking availability logic — owned by Appointment Manager
- Task creation, assignment, or completion — owned by Task Manager
- Message routing or SLA ownership — owned by Communication Hub
- Staff availability decisions — owned by Rota Manager
- Business analytics or reporting — owned by Smart Dashboards
- Calculating churn risk predictions or loyalty signals — owned by Loyalty Insights (Aiden surfaces signals produced by that module only)
- Content governance for the approved knowledge base — owned by Knowledge, Training & Learning (Aiden consumes that module's APIs and does not duplicate or override its governance)
- Validating compliance requirements or interpreting policy — owned by Knowledge, Training & Learning; Aiden MUST NOT infer compliance status or policy meaning even when surfacing knowledge content

---

3. Core Objects (Normative)

3.1 AIContext (Canonical Artefact)

An AIContext is a governed runtime artefact representing the resolved state of a user's session at the moment Aiden is invoked.

Minimum required fields:

- `UserId`
- `Role` (if staff)
- `Channel` — one of `PatientApp | StaffApp | Web | Phone`
- `ScreenContext`
- `EnabledModules[]`

Access Manager enforcement context: Every AIContext instantiation MUST include a resolved permission snapshot drawn from Access Manager at the moment of invocation. This snapshot governs which data, actions, and guidance surfaces are available for the duration of that context. Aiden MUST NOT cache or reuse a permission snapshot across sessions or context resets. The permission snapshot does not become a persisted

field on AIContext but MUST be used to validate every AIAction and every data read before any content is presented to the user.

3.2 AIIntent

An AIIntent is the structured representation of what the user is attempting to achieve, as resolved by Aiden.

Fields:

- `IntentType`
- `ConfidenceScore`
- `RequiredEntities`
- `SupportedActions[]`

3.3 AIAction

An AIAction is a candidate action Aiden surfaces to a user or staff member for explicit human confirmation.

Fields:

- `ActionType`
- `TargetModule`
- `RequiresConfirmation` (bool)

Task Manager constraint: When `TargetModule` is `TaskManager`, `ActionType` MUST NOT include `Create`, `Assign`, or `Complete`. Aiden's only permitted action types toward Task Manager are suggestion and surfacing (e.g. `Suggest`). Any task creation, assignment, or completion requires explicit staff action outside of Aiden's initiation. Task Manager records AI-originated suggestions with an `Origin` value of `AI-Suggested`, which is distinct from all other creation sources.

Campaign Manager constraint: When `TargetModule` is `CampaignManager`, Aiden MUST NOT set or imply an `ActionType` that transitions a campaign from `Draft` to `Active` or any equivalent live state. All AI-generated campaign drafts MUST remain in `Draft` state; human activation is the mandatory gate before any campaign reaches active delivery. This constraint applies regardless of how the suggestion originates (staff prompt, proactive signal, or outreach support action).

3.4 EscalationEvent

An EscalationEvent is emitted when Aiden cannot resolve a user need within its permitted boundaries and must hand off to a human.

Fields:

- `Reason`
- `ConfidenceScore`
- `TargetRole`
- `Timestamp`

3.5 AuditEvent (Canonical Artefact)

An AuditEvent is an immutable, inspection-ready record of every AI-initiated or AI-assisted interaction.

Minimum required fields:

- `EventType`
- `Actor` — always `AI` for Aiden-emitted events
- `Context`
- `Timestamp`
- `DeviceId` — aligns with the platform-wide canonical `SecurityEvent` model defined in the Security and Privacy spec, enabling device-level correlation during inspection
- `ai_involved` — bool, always `true` for AI-emitted events

AuditEvent records are immutable from the moment of creation and MUST NOT be modified or deleted by any process, including Aiden itself.

SecurityEvent stream integration: Aiden's AuditEvents MUST be emitted to the platform-wide audit log maintained by the Audit & Compliance module, which consolidates them into the canonical SecurityEvent stream as defined in the Security and Privacy spec. Aiden MUST NOT maintain a separate, siloed audit store. The `DeviceId` field on every AuditEvent is the join key that enables device-level correlation between Aiden's AI activity records and the broader SecurityEvent stream. Aiden does not own or duplicate SecurityEvent enumeration — it emits well-formed AuditEvents and the platform audit layer is responsible for classifying them within the SecurityEvent taxonomy.

3.6 CallThread DetectedIntents Contract

When Aiden is active during a Communication Hub CallThread interaction, it MUST populate the `DetectedIntents` field on the relevant CallThread sub-type as defined by Communication Hub's integration contract. The following rules apply:

- Aiden MUST write `DetectedIntents` as a structured array of resolved AIIntent snapshots, each carrying `IntentType`, `ConfidenceScore`, and `Timestamp`.
- Aiden MUST write `DetectedIntents` at the point each intent is resolved during the call, not retrospectively at call end.
- Aiden MUST NOT overwrite a previously written `DetectedIntent` entry; new intents are appended.
- If no intent can be resolved with sufficient confidence, Aiden MUST write a single entry with `IntentType: Unresolved` and the current `ConfidenceScore`, rather than leaving the field empty.
- All writes to `DetectedIntents` MUST be accompanied by a corresponding AuditEvent carrying `ai_involved: true` and the relevant `DeviceId`.
- Validation of the `DetectedIntents` payload format (field names, required values, and allowed `IntentType` values) is governed by Communication Hub's published integration schema; Aiden MUST consume and conform to that schema.

3.7 AIIntent / AIAction State Machine (Authoritative)

States:

- `Pending` — intent resolved, awaiting user or staff action
- `Presented` — action surfaced to user; awaiting confirmation
- `Confirmed` — human has explicitly approved the action
- `Escalated` — confidence threshold not met or action not supported; handed to a human
- `Abandoned` — user exited without completing

Rules:

- State transitions are auditable and time-stamped.
- An AIAction MUST NOT proceed to `Confirmed` without an explicit human trigger — Aiden cannot self-confirm.
- Once `Escalated`, an intent MUST NOT revert to `Pending` or `Presented`.
- All transitions emit an AuditEvent.

---

4. Core Capability Areas

4.1 Action-First Principle (Non-Negotiable)

If an action can be completed inside Primoro, Aiden MUST guide the user to that action rather than explain it abstractly. Buttons and deep-links are preferred over free-text instructions.

The module MUST:

- Present actionable UI controls (buttons, deep-links) wherever a supported action exists
- Default to guidance such as "You can reschedule this appointment now — just choose a new time" rather than abstract explanation
- Guide users to open incomplete forms directly (e.g. "Your medical history is due — I can open it for you")

The module MAY:

- Provide plain-English explanation when no in-app action is available and explanation aids user understanding

The module MUST NOT:

- Replace an available in-app action with a free-text description of how to do it

4.2 Screen-Aware & Contextual Guidance (Authoritative)

Aiden adapts its guidance based on current screen, user role, enabled modules, and CORE calendar context (what is due or upcoming).

The module MUST:

- Resolve an AIContext at the point of each interaction, incorporating screen, role, channel, and enabled modules
- Surface only guidance and actions relevant to the resolved context
- Surface calendar-aware signals (due, overdue, missing) drawn from the CORE calendar

The module MAY:

- Suppress the Aiden surface entirely when no clarity or guidance value is offered in the current context
- Surface recommended alternative leave dates to staff when HR & People Manager is enabled and the relevant signal is available — this is a read-only advisory surface only; Aiden MUST NOT perform leave approval, conflict detection, risk classification, or any other action that alters HR policy or approval workflows

The module MUST NOT:

- Present actions or data outside the user's resolved permissions

- Bypass scheduling, ownership, or SLA rules governed by other modules

4.3 Human-Controlled AI (Non-Negotiable)

Aiden MUST require explicit human action before any suggestion is executed. Ownership always remains with staff or patients.

The module MUST:

- Mark every surfaced action with `RequiresConfirmation: true` unless the action is purely navigational (e.g. opening a screen)
- Ensure all actions follow existing module governance
- Enforce the Campaign Manager draft-gate: any AI-generated campaign content or outreach suggestion MUST remain in `Draft` state and MUST require explicit human activation before it can be delivered; Aiden MUST NOT initiate or imply activation

The module MUST NOT:

- Auto-send messages or take autonomous actions
- Widen access beyond what Access Manager permits
- Become authoritative — suggestions are proposals, not decisions

4.4 Smart Dashboard Signal Handling (Authoritative)

When surfacing overdue or at-risk work to staff, Aiden consumes structured alert signals from Smart Dashboards. Each signal carries a status value — `OnTrack`, `Attention`, or `Overdue` — together with a severity level, escalation state, and a linked entity ID.

The module MUST:

- Interpret these semantics consistently: an `Attention` signal warrants a prompt; an `Overdue` signal warrants a prioritised, visible call to action
- Present and contextualise signals exactly as assigned by Smart Dashboards

The module MUST NOT:

- Duplicate or contradict the status that Smart Dashboards has already assigned to an entity
- Re-evaluate alert signals independently

Smart Dashboards audit-event consumption: Smart Dashboards emits audit events for dashboard access and action launches. When Aiden is active on a surface that includes a Smart Dashboards view, it MAY consume these audit events to provide contextual guidance or inline suggestions that are relevant to the entity the staff member is currently viewing. Aiden MUST NOT re-interpret or modify the content of those audit events; it uses them only as a contextual signal to determine which guidance is most relevant at that moment. Any guidance surfaced on the back of a Smart Dashboards audit event MUST itself be logged in a corresponding AuditEvent with `ai_involved: true`.

4.5 Clinical Safety (Non-Negotiable)

The module MUST:

- Respond to any symptom-based prompt with the safe fallback: "I can't assess symptoms, but if you're worried or in pain it's best to contact the practice."

- Escalate immediately on any emergency indicator

The module MUST NOT:

- Diagnose or comment on clinical outcomes under any circumstances
- Provide advice that could be construed as clinical guidance
- Validate compliance requirements or interpret policy on behalf of any module, including Knowledge, Training & Learning; compliance validation and policy interpretation are owned by Knowledge, Training & Learning and MUST NOT be inferred or approximated by Aiden even when surfacing knowledge content

4.6 Escalation (Strict)

Escalation to a human occurs only when all of the following apply:

- The action is not supported in-app
- The question cannot be answered from approved data
- The user explicitly asks for a person, or Aiden's confidence score is below the configured threshold

Default fallback wording: "I can't do that directly, but here's what I can help with…"

All escalations emit an EscalationEvent and an AuditEvent.

4.7 Aftercare Guidance Boundaries (Authoritative)

When Aftercare Manager is enabled, Aiden MAY answer approved aftercare questions on behalf of patients, surface aftercare instruction content drawn from Aftercare Manager's approved content store, and monitor sentiment indicators within the aftercare context. The following boundaries apply:

The module MUST:

- Source all aftercare content exclusively from Aftercare Manager's approved content API; Aiden MUST NOT generate or paraphrase aftercare instructions independently
- Escalate to Aftercare Manager's defined escalation threshold when a patient response or sentiment indicator meets the criteria specified by Aftercare Manager (e.g. flagged symptom language or distress signals); the EscalationEvent emitted in this case MUST carry a `TargetRole` value aligned with the role defined in Aftercare Manager's escalation rules
- Log all aftercare-related AI interactions with `ai_involved: true` and the relevant `DeviceId`

The module MUST NOT:

- Diagnose, assess, or comment on post-treatment clinical outcomes under any circumstances
- Override or bypass Aftercare Manager's escalation thresholds or flag criteria
- Present aftercare content that has not been approved within Aftercare Manager's content governance workflow

Aiden does not own aftercare escalation logic — it enforces escalation thresholds as defined and published by Aftercare Manager. Aftercare Manager is the authoritative source for what constitutes an escalation-worthy signal in this context.

4.8 Loyalty Insights Signal Handling (Authoritative)

When Loyalty Insights is enabled, Aiden MAY surface at-risk patient signals, loyalty status indicators, and proportionate action suggestions produced by Loyalty Insights to relevant staff roles.

The module MUST:

- Consume loyalty signals via the Loyalty Insights inbound API and present them as-is, without modification or independent re-scoring
- Handle null or unavailable loyalty signal states gracefully: if Loyalty Insights is unavailable or returns no signal for a given patient, Aiden MUST NOT infer a loyalty or churn status and MUST suppress the loyalty signal surface entirely for that patient until a valid signal is available
- Surface loyalty signals only to staff roles that have permission to view loyalty data as resolved from Access Manager
- Log all loyalty signal surfacing events with `ai_involved: true`

The module MUST NOT:

- Calculate, score, or predict loyalty status or churn risk independently of Loyalty Insights
- Surface a loyalty signal that has not been produced and published by Loyalty Insights

---

5. Delivery Surfaces & Access (Authoritative)

5.1 Web Portal

Aiden appears in the staff web portal providing role-aware how-to guidance, operational clarity ("what's due / missing"), suggested next actions during patient conversations, and internal Q&A using practice-approved knowledge. Aiden MUST NOT send messages or take actions toward patients without explicit staff intent.

5.2 Tablet App (Staff App Mode)

Aiden operates in Staff App Mode on the in-practice tablet with equivalent capabilities to the web portal: role-aware guidance, at-risk work surfacing, and suggested next actions. The resolved AIContext in this surface MUST include the `StaffApp` channel value so that audit records correctly identify the originating surface.

5.3 Patient Mobile App

Aiden appears in the Patient App and supports:

- Managing appointments (view / book / reschedule / cancel within rules)
- Explaining appointment types and preparation
- Explaining treatment plans and bundled pricing (non-clinical)
- Highlighting missing forms or actions and launching completion
- Viewing balances and payment options (where Integrated Payments is enabled)
- Answering approved FAQs
- Providing aftercare guidance (non-diagnostic)

Tone, safety rules, and wording are governed globally.

5.4 Website Chat (When Treatment Pipeline Enabled)

When enabled, Aiden answers common pre-treatment FAQs, collects enquiry details, creates lead callbacks (calendar-backed), and supports permitted appointment booking. Capabilities appear only when the relevant module is enabled.

Treatment Pipeline constraints in this channel (authoritative):

- Aiden MUST create or link to exactly one Lead record per intake interaction; duplicate Lead creation is prohibited.
- Aiden MUST NOT create Opportunity records — Opportunities are created only by authorised staff following triage.
- Aiden MUST NOT set or infer `BookingEligibility` — that field is controlled exclusively by the Treatment Pipeline triage process.
- Aiden MUST NOT advance pipeline stages on its own initiative.
- Aiden MUST NOT book appointments through the Treatment Pipeline channel unless `BookingEligibility` has already been explicitly set by an authorised staff member.

Aiden's role in this channel is intake and signposting only.

5.5 AI Phone Assistant (Optional Module)

When the AI Phone Assistant module is enabled, Aiden acts as a voice receptionist, handles repetitive calls (booking, FAQs, recalls), operates 24/7, escalates to humans when required, and books or records follow-ups using the same rules as all other channels. Voice is an interface, not a different logic layer.

5.6 Engagement Signals

Aiden emits the following signals for staff visibility and governance review:

- Escalation frequency by channel and intent type
- Unhandled or unanswered intents
- AI-initiated actions and their confirmation/abandonment outcomes
- Channel performance metrics

These signals are emitted as telemetry and MUST NOT contain patient-identifiable data beyond what Access Manager permits for the consuming role.

---

6. Integration Contracts

6.1 Inbound (this module consumes from)

From module	What	Contract
Appointment Manager	Appointment availability and booking rules	API (sync)
Smart Dashboards	Alert signals (`OnTrack` / `Attention` / `Overdue`) with severity, escalation state, and linked entity ID	API (sync)
Digital Forms	Form completion status and form launch targets	API (sync)

Access Manager	RBAC permissions, role context, verified patient identity	API (sync)
Knowledge, Training & Learning	Approved knowledge base for FAQ and staff Q&A; course and training pathway recommendations (when enabled)	API (sync)
Loyalty Insights	Churn risk predictions, leading indicators, proportionate action suggestions	API (sync)
Treatment Pipeline	Lead records, `BookingEligibility` status, pipeline stage	API (sync)
Integrated Payments	Balance and payment option data (when enabled)	API (sync)
CORE Calendar	Task, callback, and appointment state for calendar-aware surfacing	API (sync)
Aftercare Manager	Approved aftercare content, escalation thresholds, and sentiment flag criteria	API (sync)
HR & People Manager	Recommended alternative leave date signals (advisory, read-only; when enabled)	API (sync)
Communication Hub	CallThread context and `DetectedIntents` schema for AI-assisted call interactions	API (sync)

6.2 Outbound (this module emits to)

To module	What	Contract
Task Manager	`AI-Suggested` suggestion records (Suggest only — no Create / Assign / Complete)	event
Appointment Manager	Booking or reschedule intent (for human confirmation)	event

Communication Hub	Escalation triggers for human follow-up; `DetectedIntents` payload written to CallThread sub-type during active call interactions	event / API (sync)
Audit & Compliance	AuditEvents with `ai_involved: true` and `DeviceId`, emitted to the platform-wide SecurityEvent stream	event (immutable)
Treatment Pipeline	New Lead record per intake interaction (website channel only)	API (sync)

6.3 Communication Hub — AIIntent Thread Contract

When Aiden surfaces suggestions within a Communication Hub thread (for example, a staff member viewing an active patient conversation), the following contract applies:

- AIIntent objects surfaced in thread context MUST include `IntentType`, `ConfidenceScore`, `RequiredEntities`, and `SupportedActions[]` as defined in §3.2.
- The AIIntent payload MUST be persisted against the thread record by Communication Hub; Aiden emits the payload and Communication Hub owns thread-level persistence.
- Dismissal of a suggestion by a staff member MUST be captured as an `Abandoned` state transition on the AIIntent (see §3.7) and MUST emit a corresponding AuditEvent logged to Audit & Compliance.
- Confirmation of a suggestion MUST follow the standard `Confirmed` state transition and MUST NOT proceed without explicit staff action.
- Aiden MUST NOT write directly to the Communication Hub thread record; all thread persistence is handled by Communication Hub on receipt of Aiden's emitted payload.
- The payload schema (field names, types, and validation rules) for AIIntent objects in thread context is governed by Communication Hub's published integration schema; Aiden MUST conform to that schema.

6.4 Knowledge, Training & Learning — Course Recommendation Contract

When Knowledge, Training & Learning is enabled and course recommendations are configured, Aiden MAY surface learning course or training pathway suggestions to staff as part of action-first guidance. The following boundaries apply:

- Aiden MUST source all course recommendations exclusively from Knowledge, Training & Learning's recommendation API; it MUST NOT generate, rank, or prioritise courses independently.
- Course recommendations surfaced by Aiden MUST be visually distinguished as AI-generated suggestions and MUST be presented as secondary to any assigned work or mandatory training already surfaced on screen.
- Aiden MUST NOT enrol a staff member in a course, mark a course as complete, or alter any training record; it surfaces the recommendation and the staff member takes action.
- The `ActionType` for a course recommendation MUST be `Suggest`; `RequiresConfirmation` MUST be `true`.

- Aiden MUST NOT validate or infer compliance status based on course completion data; compliance tracking is owned by Knowledge, Training & Learning.

6.5 PMS Boundary

Aiden does not interact directly with the PMS. All scheduling, record, and clinical data flows pass through the relevant Primoro modules (Appointment Manager, Digital Forms, etc.), which own the PMS integration boundary. Aiden consumes Primoro module APIs — it does not read from or write to the PMS directly.

---

7. AI Boundaries (Non-Negotiable)

AI MAY:

- Explain Primoro features and options to patients and staff in plain English
- Guide users to complete supported actions inside Primoro via buttons and deep-links
- Surface calendar-aware signals (due, overdue, missing) for human review
- Summarise at-risk or overdue work for staff, drawing from Smart Dashboards signals
- Suggest next actions to staff using practice-approved knowledge (with explicit human confirmation before any action is taken)
- Answer approved FAQs and surface relevant training content from Knowledge, Training & Learning
- Surface course and training pathway recommendations produced by Knowledge, Training & Learning (visually secondary to assigned work and clearly badged as AI-generated)
- Surface churn risk predictions and leading indicators produced by Loyalty Insights
- Highlight potential unlinked dependants and prompt staff to consider Family Profile creation
- Surface aftercare guidance drawn from Aftercare Manager's approved content store (non-diagnostic)
- Surface recommended alternative leave dates to staff on an advisory, read-only basis (when HR & People Manager is enabled)
- Populate `DetectedIntents` on Communication Hub CallThread records during active call interactions, per the contract in §3.6

AI MAY NOT:

- Auto-decide on any policy-bound action or take autonomous action of any kind
- Diagnose, comment on clinical outcomes, or provide clinical advice
- Bypass governance, audit, or access checks
- Make commitments on behalf of the practice
- Replace required clinical judgement
- Widen access beyond what Access Manager permits
- Create Family Profile links, assign permissions, or override consent
- Create Opportunity records in the Treatment Pipeline
- Set or infer `BookingEligibility`
- Advance Treatment Pipeline stages on its own initiative
- Create, assign, or complete tasks in Task Manager
- Re-evaluate or contradict alert signal statuses assigned by Smart Dashboards

- Duplicate or override content governance in Knowledge, Training & Learning
- Validate compliance requirements or interpret policy
- Calculate loyalty or churn signals independently of Loyalty Insights
- Transition a Campaign Manager campaign from `Draft` to `Active` or any live state
- Enrol staff in courses, mark courses complete, or alter any training record in Knowledge, Training & Learning
- Perform leave approval, conflict detection, or risk classification in HR & People Manager
- Generate or paraphrase aftercare instructions independently of Aftercare Manager's approved content
- Override or bypass Aftercare Manager's escalation thresholds or flag criteria
- Write directly to Communication Hub thread records (thread persistence is owned by Communication Hub)

---

8. Audit & Compliance

The system MUST log:

- All AI-initiated prompts, with channel, role context, and screen context
- All suggested actions surfaced to users, including which were confirmed, abandoned, or escalated
- All escalation events, with reason, confidence score, target role, and timestamp
- All actions launched via Aiden (deep-link or button), with the target module and action type
- All AI involvement in any cross-module event, flagged with `ai_involved: true`
- All unanswered or escalated intents (telemetry, for continuous improvement)
- All `DetectedIntents` written to Communication Hub CallThread records
- All loyalty signal surfacing events
- All aftercare content surfacing events and any escalation threshold triggers arising from them
- All course recommendation surfacing events (Knowledge, Training & Learning channel)

Every AuditEvent emitted by Aiden MUST carry:

- `ai_involved: true`
- `DeviceId` — aligned with the platform-wide canonical `SecurityEvent` model in the Security and Privacy spec, enabling device-level correlation during inspection

Audit logs MUST be immutable and exportable for inspection. All AuditEvents are emitted to the platform-wide audit log maintained by the Audit & Compliance module and consolidated into the SecurityEvent stream; Aiden does not maintain a separate audit store. Learning and telemetry data captured under §7.2 MUST NOT bypass privacy or access controls.

---

9. Access Control

Aiden respects Access Manager RBAC at all times. Access rules:

- **Patient users** — may access only their own data within the Patient App surface; identity MUST be verified before any patient-bound data is presented

- **Staff users** — may access data and actions appropriate to their assigned role; Aiden MUST resolve the role from Access Manager at each AIContext instantiation
- **No cross-role leakage** — Aiden MUST NOT surface data or actions outside the resolved permissions for the active user
- **AI never widens access** — Aiden MUST NOT grant, infer, or assume permissions not explicitly assigned by Access Manager

MFA requirements for sensitive operations (e.g. accessing patient records in certain contexts) are governed by Access Manager policy; Aiden MUST honour any MFA gate enforced by Access Manager before presenting protected data or actions.

---

10. Integration Summary

- **Appointment Manager** — inbound booking rules and availability; outbound booking/reschedule intent (human-confirmed)
- **Task Manager** — outbound `AI-Suggested` suggestion records only; no Create / Assign / Complete
- **Communication Hub** — outbound escalation triggers for human follow-up; outbound `DetectedIntents` payload written to CallThread sub-type during active call interactions; AIIntent suggestion payloads surfaced within thread context per the contract in §6.3
- **Digital Forms** — inbound form completion status and form launch targets
- **Aftercare Manager** — inbound approved aftercare content, escalation thresholds, and sentiment flag criteria for non-diagnostic patient surfacing
- **Document Hub** — inbound document context for staff guidance surfaces
- **Access Manager** — inbound RBAC, role context, and patient identity verification; governs all access decisions and provides the permission snapshot resolved at every AIContext instantiation
- **Smart Dashboards** — inbound alert signals consumed and contextualised; never re-evaluated; audit events from dashboard access and action launches MAY be consumed as contextual signals for inline guidance
- **Treatment Pipeline** — inbound Lead records and `BookingEligibility`; outbound new Lead creation (website channel only); strict intake-only constraints apply
- **Integrated Payments** — inbound balance and payment option data (when enabled)
- **Loyalty Insights** — inbound churn risk predictions and leading indicators (when enabled); Aiden surfaces, does not calculate; null/unavailable states suppress the loyalty signal surface
- **Knowledge, Training & Learning** — inbound approved knowledge base for FAQ and staff Q&A; inbound course and training pathway recommendations (when enabled); Aiden consumes APIs, does not duplicate or override content governance, and MUST NOT validate compliance or interpret policy
- **HR & People Manager** — inbound recommended alternative leave date signals (advisory, read-only; when enabled); Aiden does not perform leave approval, conflict detection, or risk classification
- **Audit & Compliance** — outbound immutable AuditEvents with `ai_involved` flag and `DeviceId`, emitted to the platform-wide SecurityEvent stream
- **AI Phone Assistant** — optional voice channel extension; same logic layer as all other channels
- **AI Quality Monitor** — optional extension for quality and compliance insights
- **AI Guardian** — optional extension for anomaly and risk detection; form relevance suggestions as defined in the Digital Forms extension contract

- **Campaign Manager** — optional extension for outreach support; all AI-generated campaign content MUST remain in `Draft` state pending explicit human activation
- **CORE Calendar** — inbound task, callback, and appointment state for calendar-aware surfacing

---

11. Explicit Non-Goals

- **Replacing reception or admin teams** — Aiden reduces repetitive interruptions but does not substitute for human staff; no module currently scoped to own autonomous reception workflows
- **Performing autonomous actions** — any future autonomous action capability would require a new governed module with its own safety specification
- **Overriding platform rules or policies** — if added, policy override tooling would be owned by a governance-layer module
- **Providing analytics or reporting** — owned by Smart Dashboards
- **Acting as a clinical system** — clinical records and clinical decision support are explicitly out of scope for the Primoro platform

---

12. Versioning & Governance

This specification is owned by: the AI Assistant module owner.

Changes to this spec require:

- Review by the MVP module owner
- Impact analysis across declared related modules (see /propose)
- Version bump (patch / minor / major) depending on scope of change

All future changes must preserve:

- Action-first guidance
- Human control over all AI suggestions
- Clinical safety boundaries
- CORE calendar governance
- Modular enablement (capabilities appear only when relevant modules are enabled)

---

13. Build Contract (Engineering & QA)

13.1 Canonical Data Model

AIContext

Field	Type	Notes
`UserId`	UUID	Verified identity from Access Manager

`Role`	Enum	Staff role from Access Manager; null for patients
`Channel`	Enum	`PatientApp | StaffApp | Web | Phone`
`ScreenContext`	String	Current screen identifier
`EnabledModules`	String[]	List of enabled module slugs for this practice

AIIntent

Field	Type	Notes
`IntentType`	String	Resolved intent category
`ConfidenceScore`	Float	0.0–1.0
`RequiredEntities`	String[]	Entities needed to fulfil intent
`SupportedActions`	String[]	Actions available given current context

AIAction

Field	Type	Notes
`ActionType`	String	MUST NOT be `Create`, `Assign`, or `Complete` when `TargetModule` is `TaskManager`; MUST NOT imply activation when `TargetModule` is `CampaignManager`
`TargetModule`	String	Canonical module slug
`RequiresConfirmation`	Bool	MUST be `true` for all non-navigational actions

EscalationEvent

Field	Type	Notes
`Reason`	String	Why escalation was triggered
`ConfidenceScore`	Float	Aiden's confidence at point of escalation

`TargetRole`	String	Role to receive escalation; MUST align with Aftercare Manager's defined escalation role when triggered in an aftercare context
`Timestamp`	DateTime	UTC

AuditEvent

Field	Type	Notes
`EventType`	String	Category of event
`Actor`	String	Always `AI` for Aiden-emitted events
`Context`	Object	Channel, role, screen, intent snapshot
`Timestamp`	DateTime	UTC
`DeviceId`	UUID	Aligns with platform `SecurityEvent` model; join key for SecurityEvent stream correlation
`ai_involved`	Bool	Always `true` for Aiden-emitted events

13.2 Core Behaviour Rules

1. Aiden MUST resolve an AIContext at the start of every interaction; no guidance may be presented without a resolved context.
2. Aiden MUST NOT present any action with `RequiresConfirmation: false` unless the action is purely navigational (opening a screen).
3. Aiden MUST NOT set `ActionType` to `Create`, `Assign`, or `Complete` when `TargetModule` is `TaskManager`.
4. Aiden MUST escalate when: the action is not supported in-app AND the question cannot be answered from approved data AND (the user asks for a human OR the `ConfidenceScore` is below the configured escalation threshold).
5. Aiden MUST respond to any symptom-based prompt with the mandated safe fallback message and MUST NOT provide clinical commentary.
6. Aiden MUST emit an EscalationEvent and an AuditEvent for every escalation.
7. Every AuditEvent MUST carry `ai_involved: true` and a valid `DeviceId`, and MUST be emitted to the platform-wide audit log (Audit & Compliance module / SecurityEvent stream).
8. Aiden MUST NOT create more than one Lead record per intake interaction in the website channel.

9. Aiden MUST NOT create Opportunity records or set/infer `BookingEligibility` in the Treatment Pipeline channel.

10. Aiden MUST present Smart Dashboards alert signals with the status already assigned (`OnTrack` / `Attention` / `Overdue`) and MUST NOT re-evaluate those statuses.

11. Aiden MUST NOT surface data or actions outside the permissions resolved from Access Manager for the active user.

12. All state transitions on AIIntent and AIAction objects MUST be time-stamped and produce an AuditEvent.

13. Aiden MUST NOT transition a Campaign Manager campaign from `Draft` to any live state; all AI-generated campaign drafts require explicit human activation.

14. When populating `DetectedIntents` on a Communication Hub CallThread, Aiden MUST write each resolved intent as it occurs (not retrospectively), MUST NOT overwrite existing entries, and MUST write an `Unresolved` entry rather than leave the field empty when no intent can be resolved with sufficient confidence.

15. When Loyalty Insights is unavailable or returns no signal, Aiden MUST suppress the loyalty signal surface for that patient and MUST NOT infer a loyalty or churn status.

16. Aiden MUST source all aftercare content from Aftercare Manager's approved content API and MUST escalate when a patient response meets Aftercare Manager's defined escalation thresholds.

17. Aiden MUST NOT validate compliance requirements or interpret policy, including when surfacing content from Knowledge, Training & Learning.

18. When surfacing course recommendations from Knowledge, Training & Learning, Aiden MUST badge them as AI-generated, present them as secondary to assigned work, and MUST use `ActionType: Suggest` with `RequiresConfirmation: true`.

19. When surfacing HR leave date suggestions, Aiden MUST treat these as read-only advisory signals only; it MUST NOT perform or initiate leave approval, conflict detection, or risk classification.

13.3 Configuration Surfaces

Practice-level settings (Admin Control Plane):

- Enabled channels (`PatientApp`, `StaffApp`, `Web`, `Phone`)
- Supported intent types
- Escalation confidence threshold
- FAQ knowledge sources
- Voice enablement (AI Phone Assistant)
- Language and tone variants

Per-user preferences (Access Manager):

- Role assignment (governs which guidance surfaces are visible)

All configuration changes MUST be logged with actor and timestamp in the platform audit trail.

13.4 Filtering & Views

Admins and governance users can view:

- Escalation frequency by channel and intent type
- Unhandled or unanswered intents
- AI-initiated actions and their confirmation / abandonment outcomes

- Channel performance metrics

All views in this surface MUST enforce Access Manager role permissions; governance users MUST NOT see patient-identifiable detail beyond their permitted access level.

13.5 Module Extension Map

Optional extensions that extend Aiden without breaking this contract:

Extension Module	Capability Added
AI Phone Assistant	Voice channel (`Phone` value in `Channel` enum)
AI Quality Monitor	Quality and compliance insights surface
AI Guardian	Anomaly and risk detection; form relevance suggestions (Digital Forms extension contract)
Campaign Manager	Outreach support actions (Draft-only; human activation gate enforced)

Each extension MUST conform to the same AIContext, AIAction, EscalationEvent, and AuditEvent contracts defined in §13.1. Extensions MUST NOT introduce new `ActionType` values that bypass the `RequiresConfirmation` requirement.

13.6 Acceptance Criteria

The build of AI Assistant (Aiden) is complete when:

- [ ] All canonical objects (AIContext, AIIntent, AIAction, EscalationEvent, AuditEvent) can be created and read through the API
- [ ] State machine transitions on AIIntent / AIAction enforce all rules in §3.7
- [ ] All integrations in §6 are wired
- [ ] AI boundaries in §7 are enforced (negative tests pass for all MUST NOT items)
- [ ] Audit log captures every event in §8, with `ai_involved: true` and `DeviceId` on every AI-emitted event; all events emitted to platform-wide SecurityEvent stream via Audit & Compliance module
- [ ] Access control is enforced per §9; no cross-role leakage test passes; Access Manager permission snapshot is resolved at every AIContext instantiation
- [ ] Treatment Pipeline website-channel constraints enforced (no Opportunity creation, no `BookingEligibility` inference, exactly one Lead per intake interaction)
- [ ] Task Manager action-type prohibitions enforced (no `Create`, `Assign`, or `Complete` via Aiden)
- [ ] Smart Dashboards alert signal semantics respected and not contradicted
- [ ] Clinical safety fallback triggered correctly on symptom-based prompts
- [ ] Escalation rules enforced: EscalationEvent and AuditEvent emitted on every escalation
- [ ] Campaign Manager draft-gate enforced: no AI-generated campaign content transitions to Active state
- [ ] CallThread `DetectedIntents` populated correctly per §3.6: incremental writes, no overwrites, `Unresolved` entry on low-confidence intents

- [ ] Communication Hub AIIntent thread contract enforced per §6.3: dismissal captured as `Abandoned`, confirmation requires explicit staff action
- [ ] Loyalty Insights null/unavailable state handled: loyalty signal surface suppressed when no valid signal available
- [ ] Aftercare content sourced exclusively from Aftercare Manager approved content API; escalation threshold triggers produce correct EscalationEvent with aligned `TargetRole`
- [ ] Compliance validation and policy interpretation excluded from all Aiden responses
- [ ] Knowledge course recommendations badged as AI-generated, presented as secondary to assigned work, and emitted with `ActionType: Suggest` and `RequiresConfirmation: true`
- [ ] HR leave date suggestions surfaced as read-only advisory signals only; no leave approval, conflict detection, or risk classification initiated
- [ ] All non-functional requirements in §14 are met

---

14. Non-Functional Requirements

- **Performance:** Aiden MUST deliver conversational responses with low latency; all module API calls MUST be non-blocking so that a slow upstream module does not stall the Aiden surface. Target response latency and specific throughput numbers are not captured in the original and must be defined by engineering before build-contract lock.
- **Reliability:** Aiden MUST degrade gracefully on upstream API failure — if a module API is unavailable, Aiden MUST fall back to the escalation flow rather than presenting incorrect or stale data. Safe fallback to human workflows is the default degraded state.
- **Scalability:** Aiden operates as a shared guidance layer across all practices on the platform; the design MUST support multi-tenant isolation so that one practice's configuration, knowledge sources, and telemetry do not bleed into another's.
- **Security:** All data exchanged between Aiden and upstream modules MUST be encrypted in transit. No patient-identifiable data MAY be persisted by Aiden beyond what is required to serve the active session and what is permitted by data governance policy.
- **Privacy:** Aiden MUST honour Access Manager-governed data permissions. Telemetry and learning data captured under §7 MUST NOT contain patient-identifiable information beyond what is permitted for the consuming role. Data retention policy for AuditEvent records is governed by the Audit & Compliance module; Aiden does not own retention decisions.
- **Observability:** Aiden MUST export escalation frequency, unhandled intent rates, action confirmation rates, and channel performance metrics. All AuditEvents MUST be emitted to the Audit & Compliance module in real time. Specific metrics endpoints, trace formats, and alerting thresholds are not captured in the original and must be defined by engineering before build-contract lock.
- **Accessibility:** Aiden's conversational surfaces in the Patient App and Web Portal MUST meet the accessibility standards applicable to the platform; specific conformance level is not captured in the original and must be defined before build-contract lock.

---

15. Open Questions

> Outstanding decisions before this spec can be promoted from `draft` to `published`.

1. **Escalation confidence threshold:** The spec states that escalation occurs when "AI confidence is low" and that the threshold is configurable, but no default value or range is defined. What is the default confidence threshold below which Aiden escalates, and who sets the floor?

2. **Performance targets:** Low-latency responses and non-blocking API calls are required, but no specific latency or throughput figures are captured. What are the engineering-agreed targets for P50, P95, and P99 response times?

3. **Voice channel (AI Phone Assistant) logic parity:** The spec states "voice is an interface, not a different logic layer" but does not define how AIContext, EscalationEvent, and AuditEvent are produced in a real-time voice session. Does the AI Phone Assistant module own this adaptation, or does Aiden?

4. **Telemetry retention and patient-identifiability:** Section 7.2 states learning data never bypasses privacy controls, but no retention period for telemetry records is specified. What is the retention policy for non-audit telemetry (unanswered questions, repeated intents, abandoned flows)?

5. **Accessibility conformance level:** No WCAG or platform accessibility conformance target is specified for Aiden's UI surfaces. What is the required conformance level?

6. **Tone and language variants:** Admin configuration supports language and tone variants, but no supported locales or approved variants are listed. Which languages and tone profiles are in scope for MVP?

7. **AI Guardian form relevance contract:** The Module Extension Map references "the Digital Forms extension contract" for AI Guardian's form relevance suggestions, but this contract is not defined in this spec. Where is that contract authoritative, and is it owned by AI Guardian or Digital Forms?

8. **CallThread DetectedIntents schema ownership:** §3.6 states that the `DetectedIntents` payload schema is governed by Communication Hub's published integration schema. Where is that schema published, and what is the versioning and change-notification process when Communication Hub updates it?

9. **Aftercare escalation role alignment:** §4.7 requires that EscalationEvent `TargetRole` aligns with the role defined in Aftercare Manager's escalation rules. How are changes to Aftercare Manager's escalation role configuration communicated to Aiden at runtime, and is this configuration-driven or requires a deployment?

10. **HR leave date signal availability:** §4.2 permits Aiden to surface recommended alternative leave dates when HR & People Manager is enabled. What is the expected signal format and API endpoint, and under which conditions is no signal returned?