Document Hub

Doc type: technical · **Version:** v0.1 · **Status:** published · **Module slug:** document-hub
Exported: 2026-05-15 11:10 UTC · **By:** anonymous

Document Hub – Technical Specification

1. Module Purpose & Scope (Authoritative)

Document Hub is the single, secure source of truth for all practice documents in Primoro. It solves the problem of uncontrolled document duplication across systems by ensuring every document lives once, centrally, and is accessed everywhere else by secure reference — never as an uncontrolled copy. It sits at the foundation of the patient and staff journey, governing every point at which a document is created, stored, shared, acknowledged, or destroyed.

It governs:

- End-to-end document ingestion from staff uploads, patient submissions, scanner and email sources, and Primoro module outputs (forms, treatment plans, aftercare, agreements).
- Document lifecycle management from creation through approval, supersession, archiving, and controlled deletion, including retention policy enforcement and GDPR right-to-erasure.
- Secure viewing, staff-only annotation, acknowledgement tracking, and all internal, patient-initiated, and third-party sharing of documents.

It explicitly does not:

- Own business workflow logic — that is the responsibility of Task Manager and Appointment Manager.
- Perform clinical decision-making — that is the responsibility of the PMS (e.g. Dentally).
- Orchestrate messaging — that is the responsibility of Communication Hub.
- Enforce authentication or role-based access control — that is the responsibility of Access Manager.
- Surface analytics or performance dashboards — that is the responsibility of Smart Dashboards.
- Retain or govern Smart Treatment Proposal artefacts — Smart Treatment Proposals retains and governs its own document artefacts within its own secure storage boundary and those artefacts are NOT ingested into Document Hub.

2. Ownership & Responsibilities

2.1 Document Hub IS Responsible For

- Ingesting documents from all authorised sources (staff uploads, patient submissions, scanners, email, and Primoro module outputs) and storing a single master copy, encrypted, indexed, and categorised.
- Enforcing the document lifecycle state machine, maintaining immutable version history, and providing medico-legal traceability across all document events.
- Issuing secure, time-limited references to external consumers (PMS, Communication Hub, Lab Manager, patients) and revoking those references automatically when the underlying document is deleted or the recipient's permissions are withdrawn.

- Logging an immutable audit event for every document action — upload, view, download, annotation, share, revocation, acknowledgement, version change, deletion, and purge — and emitting those events to the Security & Privacy module's platform-wide audit trail in the canonical SecurityEvent shape (see §3.3 and §8).
- Enforcing real-time permission decisions issued by Access Manager, including the immediate termination of active viewing sessions when access is revoked.
- Honouring GDPR, CQC, and healthcare governance requirements, including configurable retention policies and controlled permanent erasure.

2.2 Document Hub IS NOT Responsible For

- Task creation or workflow orchestration — owned by Task Manager.
- Appointment scheduling or appointment-linked workflow — owned by Appointment Manager.
- Clinical record-keeping or clinical decision support — owned by the PMS.
- Sending notifications or messages — owned by Communication Hub (Document Hub supplies only the secure reference). Communication Hub MUST NOT store documents; Document Hub remains the sole authoritative store for all document content, and any document referenced in a Communication Hub message MUST be resolved via Document Hub's secure-reference mechanism.
- Authentication, session management, or RBAC rule definition — owned by Access Manager.
- Governance dashboards and compliance-reporting surfaces — owned by Smart Dashboards and future Governance Reporting extension.
- Storage or lifecycle management of Smart Treatment Proposal PDFs — owned by Smart Treatment Proposals.

3. Core Objects (Normative)

3.1 Document (Canonical Artefact)

A Document is a governed digital artefact representing a single versioned file (or set of related files) stored under Primoro's authority.

Minimum required fields:

- `DocumentId` — globally unique identifier
- `Category` — document type controlling RBAC and retention policy
- `PatientId` — optional; FK to patient context
- `Source` — one of: `Staff | Patient | Module | External`
- `CurrentVersionId` — FK to the active DocumentVersion
- `LifecycleState` — see §3.2
- `CreatedBy` — UserId and Role of the actor who created the record
- `CreatedAt` — timestamp of initial ingestion
- `AuditTrail` — immutable, append-only event log (see §8)

3.2 Document State Machine (Authoritative)

States:

- **Draft** — document has been ingested but not yet approved for use

- **Approved** — document is active and accessible to authorised parties
- **Superseded** — a newer version has been promoted; this version is retained but no longer current
- **Archived** — document has reached end of active use; retained for the configured retention period
- **Deleted (Pending Purge)** — soft-deleted; all access and shared links are immediately revoked; document awaits permanent erasure after retention period
- **Purged** — permanently erased after retention period has elapsed; audit record of erasure is retained

Rules:

- State transitions are auditable and time-stamped.
- Previous versions are immutable once a new version is promoted; they are retained for audit and medico-legal purposes.
- A document cannot return to Draft once it has reached Approved.
- Signed-form PDFs ingested from Digital Forms are version-locked to the signing event upon ingestion; no subsequent lifecycle operation (annotation, supersession, or re-upload) may alter the original signed artefact.
- Deletion immediately revokes all active ShareLinks and terminates any active viewing sessions for that document.
- Only users holding a role granted delete or purge permission by Access Manager may trigger the Deleted or Purged transitions.

3.3 Supporting Objects

DocumentVersion

- `VersionId`
- `DocumentId`
- `CreatedBy`
- `CreatedTimestamp`
- `FileHash`
- `StorageReference`

Annotation

- `AnnotationId`
- `DocumentVersionId`
- `CreatedBy`
- `Content`
- `Timestamp`

ShareLink

- `LinkId`
- `DocumentId`
- `CreatedBy` — `Staff | Patient`
- `ExpiryTimestamp`
- `Revoked` — boolean

AuditEvent

Document Hub's AuditEvent model extends the canonical SecurityEvent shape defined by the Security & Privacy module to ensure interoperability and a complete, immutable audit trail:

- `EventId` — globally unique, immutable; aligns with SecurityEvent.EventId
- `EventType` — enumerated: `Upload | Ingest | View | Download | Annotate | Share | Revoke | Acknowledge | VersionChange | Delete | Purge | IntegrationNotification | Escalation`; aligns with SecurityEvent.EventType
- `Actor` — structured per SecurityEvent Actor taxonomy: UserId, Role, SessionId
- `Target` — structured per SecurityEvent Target taxonomy: DocumentId, DocumentVersionId, ShareLinkId (where applicable)
- `DeviceId` — device or integration client identifier; aligns with SecurityEvent.DeviceId
- `Timestamp`

All AuditEvent records are immutable once written and MUST be retained for the full document retention period of the associated document. Document Hub MUST emit AuditEvent records covering document View, Share, Revoke, and Delete actions to the Security & Privacy module's platform-wide audit trail (see §6.2), satisfying the controlled document viewing and sharing audit obligations declared in Security & Privacy §2.1.

4. Ingestion & Creation

4.1 Staff Uploads (Authoritative)

The module MUST:

- Accept portal and staff app uploads.
- Apply immediate encryption, OCR, and indexing upon receipt.
- Categorise the document by type and patient context.
- Emit a secure reference to the PMS where relevant.

The module MUST NOT:

- Store a raw file copy in the PMS or any external system.

4.2 Patient Submissions (Authoritative)

The module MUST:

- Accept patient app and portal uploads.
- Route submitted documents for staff review before they enter an Approved state.
- Apply OCR indexing.
- Never transmit patient-submitted documents as email attachments.

4.3 Module-Generated Documents (Authoritative)

Documents generated by other Primoro modules — including Digital Forms PDFs, treatment plans, aftercare instructions, subscription agreements, care plan contracts, and other agreements — are saved automatically to Document Hub at generation time.

Digital Forms — inbound contract and immutability requirement. When a patient or staff member submits a completed form in Digital Forms, Digital Forms MUST push the signed-form PDF to Document Hub as a `Module`-sourced document at the point of generation. The push payload MUST include:

- `FormType` — the category of form completed (e.g. consent, medical history, triage)
- `PatientId` — FK to the patient who completed or for whom the form was completed
- `SignatureTimestamp` — the timestamp of the signing event, as recorded by Digital Forms
- `DelegatedSigningAttribution` — where a form was signed on behalf of a patient (e.g. by a guardian or staff member under a delegated signing rule), this MUST be included in the payload metadata; it is stored alongside the document record and included in the AuditEvent
- `SignedPdfReference` — the opaque identifier issued by Digital Forms identifying the artefact being transferred

Upon receipt, Document Hub MUST treat the ingested PDF as immutable immediately. The stored file MUST be version-locked to the signing event, and no subsequent lifecycle operation (annotation, supersession, or re-upload) may alter the original signed artefact or its `FileHash`.

Hygiene Subscriptions — signed agreement storage. When a patient subscribes to a hygiene plan, Hygiene Subscriptions MUST push the generated signed subscription agreement PDF to Document Hub as a `Module`-sourced document, practice-branded at generation time. Document Hub MUST:

- Ingest the agreement and store it as a single master copy, encrypted and indexed.
- Categorise it under a configurable subscription-agreement document category.
- Return a `SignedPdfReference` (secure reference) to Hygiene Subscriptions so the agreement can be linked to the patient record without storing a raw copy outside Document Hub.
- Treat the ingested agreement as immutable upon receipt; no subsequent lifecycle operation may alter the original signed artefact or its `FileHash`.

Care Plan Subscriptions — signed contract storage. Care Plan Subscriptions stores signed patient contracts via Document Hub. When a care plan agreement is generated, Care Plan Subscriptions MUST push the signed contract PDF to Document Hub as a `Module`-sourced document. Document Hub MUST:

- Ingest and store the contract as a single master copy, encrypted and indexed, under a configurable care-plan-contract document category.
- Return a secure reference to Care Plan Subscriptions for linkage to the patient record.
- Enforce immutability upon receipt, consistent with the signed-artefact rules above.
- Maintain a complete, immutable audit trail for all subsequent access and lifecycle events, satisfying the audit obligations declared in Care Plan Subscriptions §3.2.

Smart Treatment Proposals — explicit exclusion. Proposal PDFs generated by Smart Treatment Proposals are NOT ingested into Document Hub. Document Hub has no ownership or lifecycle responsibility over Smart Treatment Proposal artefacts. Engineers MUST NOT implement proposal PDF ingestion into Document Hub.

4.4 External Sources — Scanner and Email Ingestion (Authoritative)

The module MUST:

- Support email forwarding and scanner integrations for automated ingestion and classification.
- Flag documents that cannot be automatically associated with a patient for manual assignment, holding them in an unassigned queue pending review.

Lab-context scanner artefacts. Digital impression files received from authorised scanners (e.g. iTero) MUST be ingested, encrypted, and indexed in the same manner as other scanner artefacts. Additionally:

- Ingested impression artefacts MUST be categorised under the `LabArtefact` document category.
- Document Hub MUST expose a secure reference for each ingested artefact so that Lab Manager can attach it to the relevant Lab Case without storing a raw copy.
- Scanner-generated status events that accompany artefact delivery are treated as governed lifecycle updates and MUST be recorded in the audit log.
- Where the patient or lab case association cannot be determined automatically, the artefact MUST be held in the unassigned queue for manual review.

4.5 Referral Manager — Referral Documents and Outcome Reports (Authoritative)

Referral Manager governs the creation and receipt of referral documents and outcome reports as structured records. Document Hub provides the authoritative storage layer for these artefacts. The following rules apply:

- Referral supporting documents uploaded by staff during referral creation (e.g. X-rays, clinical notes, imaging files) MUST be ingested into Document Hub as `Staff`-sourced or `Externa l`-sourced documents (as appropriate) and linked to the relevant referral record via a secure reference held by Referral Manager.
- AI-drafted referral letter content generated by Referral Manager MUST be pushed to Document Hub as a `Module`-sourced document at the point of finalisation, before transmission to the receiving clinician or third party.
- Outcome reports and correspondence received from external referral recipients (e.g. via the referrals mailbox or email forwarding integration) MUST be ingested via the standard email ingestion pathway defined in §4.4, categorised under a configurable referral-document category, and linked to the originating referral record by Referral Manager.
- Where the patient or referral association cannot be determined automatically, the artefact MUST be held in the unassigned queue for manual review, consistent with the rules in §4.4.
- All referral document events MUST be captured in the immutable audit trail.

5. Delivery Surfaces & Access (Authoritative)

5.1 Web Portal

Staff access Document Hub via the Primoro web portal, which presents an embedded secure viewer, document list with filtering, annotation tools, acknowledgement tracking, and manager-level views of access events and external shares.

5.2 Tablet App (Staff App Mode)

In-practice tablet surfaces support document delivery, in-app patient acknowledgement capture, and staff-facing document review, operating under the same RBAC and audit controls as the web portal.

5.3 Staff App Mode — Mobile Document Access and Acknowledgement

Staff App Mode surfaces role-scoped document views sourced from Document Hub. Document Hub MUST expose the following to the Staff App Mode mobile surface:

- A filtered, RBAC-scoped list of documents relevant to the authenticated staff member's role and site, consistent with the permission model defined in §9.

- Document acknowledgement prompts for any document requiring acknowledgement by the authenticated staff member, based on the acknowledgement requirements configured per category or per document (§13.3).
- A protocol for recording acknowledgement confirmations: when a staff member confirms acknowledgement within Staff App Mode, the confirmation MUST be recorded by Document Hub as an immutable `Acknowledge` AuditEvent, capturing the actor (UserId, Role, SessionId), the document version acknowledged, site, and timestamp.
- Staff App Mode MUST NOT surface staff annotations to patients, and MUST NOT expose documents beyond the staff member's RBAC-granted category access.

Patients continue to access documents via the Patient Mobile App described in §5.4.

5.4 Patient Mobile App

Patients access documents delivered to them via in-app delivery with notifications. Patients may generate secure, time-limited links to share documents with third parties, control verification requirements, and revoke access instantly. Patients never see internal staff annotations.

5.5 Engagement Signals

Document Hub emits the following signals for staff visibility:

- Acknowledgement completion rates per document, role, and site — surfaced to managers.
- Outstanding acknowledgement gaps — available for manager filtering.
- External share activity (active links, access events, revocations) — logged and available to authorised staff.
- Integration notification events emitted to Communication Hub when document actions require patient or staff notification.

6. Integration Contracts

6.1 Inbound (this module consumes from)

From module	What	Contract
Access Manager	Permission state (role, category-level access, revocations)	Real-time; synchronous per-request; MUST NOT be cached beyond single request
Digital Forms	Signed-form PDFs at generation time, including FormType, PatientId, SignatureTimestamp, DelegatedSigningAttribution, and SignedPdfReference	Event-driven / async push
Aftercare Manager	Care documents at generation time	Event-driven / async push
Hygiene Subscriptions	Signed subscription agreement PDFs at generation time	Event-driven / async push

Care Plan Subscriptions	Signed care plan contract PDFs at generation time	Event-driven / async push
Referral Manager	Supporting documents at referral creation; AI-drafted referral letters at finalisation; outcome reports via email ingestion pathway	Event-driven / async push; email ingestion
Communication Hub	Document reference requests (secure link issuance)	Synchronous API
Lab Manager	Artefact association requests (linking secure reference to a Lab Case)	Synchronous API
PMS (e.g. Dentally)	Metadata and secure-reference consumption (read-only)	Synchronous API
Staff / Patient upload clients	Raw file upload	HTTPS multipart

6.2 Outbound (this module emits to)

To module	What	Contract
Access Manager	Permission check requests	Synchronous API
Communication Hub	Document-event notification triggers; secure time-limited references for embedding in messages	Event + synchronous reference API
Task Manager	Document-linked task triggers (e.g. unassigned document in queue)	Event / async
Lab Manager	Secure artefact reference for LabArtefact documents	Synchronous API
PMS (e.g. Dentally)	Secure metadata and Primoro document links	Synchronous API
Referral Manager	Secure document references for ingested referral artefacts	Synchronous API
Security & Privacy module	AuditEvent records extending SecurityEvent, including View, Share, Revoke, and Delete events	Async append

6.3 PMS Boundary

The PMS (e.g. Dentally) receives document metadata and secure Primoro links. The PMS MUST NOT store raw files. All access to document content via PMS-surfaced links requires Primoro authentication and is subject to Document Hub's RBAC and audit controls.

7. AI Boundaries (Non-Negotiable)

Module does not embed AI surfaces directly in the current scope.

The Module Extension Map (§13.5) identifies AI Guardian as a future extension for unusual access-pattern detection. If that extension is enabled in a future version, the following boundaries apply:

AI MAY:

- Flag unusual document access patterns for human review (AI Guardian extension only).
- Suggest document categories during ingestion for human confirmation before commit.

AI MAY NOT:

- Auto-approve, auto-share, or auto-delete any document.
- Bypass RBAC, audit logging, or retention controls.
- Make acknowledgement decisions on behalf of a user.
- Replace required human review of patient-submitted or unassigned documents.

8. Audit & Compliance

The system MUST log the following events, each as an immutable AuditEvent record (see §3.3):

- Document upload and ingest (all sources), with actor, source type, and timestamp.
- Document view and download, with actor, document version, and device identifier.
- Annotation creation, with actor, document version, and annotation content reference.
- ShareLink creation, with actor, recipient type (internal / patient / third-party), and expiry.
- ShareLink access by a third party, with verification outcome and timestamp.
- ShareLink revocation, with actor and timestamp.
- Acknowledgement events, with actor, role, site, and document version acknowledged.
- Version changes and supersession events, with actor and previous version reference.
- Deletion (soft delete), with actor, reason (where captured), and timestamp.
- Purge (permanent erasure), with confirmation of retention period elapsed.
- Permission-revocation enforcement events (session termination), with actor (system), affected user, and document.
- Integration notification events received or emitted, with source module and event type.
- Scanner-generated status events accompanying lab artefact delivery.

Audit logs MUST be immutable and append-only once written. Audit logs MUST be exportable for inspection by authorised compliance roles and MUST be retained for at least the full document retention period of the associated document.

Document Hub MUST emit AuditEvent records for document View, Share, Revoke, and Delete actions to the Security & Privacy module's platform-wide audit trail in the canonical SecurityEvent shape, as required by Security & Privacy §2.1. This emission is in addition to Document Hub's own internal immutable log; the two are not mutually exclusive.

9. Access Control

Access Manager is the authoritative source for all RBAC decisions. Document Hub MUST NOT cache permission decisions beyond the scope of a single request and MUST re-validate permissions on every secure-reference resolution.

Action	Who may perform
Upload / ingest	Roles granted upload permission for the relevant document category
View	Roles granted read permission for the relevant document category; patients for their own documents
Annotate	Roles granted annotation permission; patients MUST NOT see staff annotations
Approve / promote lifecycle state	Roles granted approve permission for the relevant category
Share internally	Roles granted share permission
Share to patient	Roles granted patient-delivery permission
Patient-initiated third-party share	Patient (for their own documents only)
Delete (soft)	Roles granted delete permission
Purge (permanent)	Roles granted purge permission
Configure retention, categories, acknowledgement requirements	Admin roles via Admin Control Plane

Document Hub MUST immediately terminate any active viewing session for a document when Access Manager revokes or modifies the viewing user's role or category-level access. The viewer MUST close or block further interaction and display an appropriate access-denied message.

MFA requirements for sensitive operations (delete, purge, and bulk-share) are governed by Access Manager policy; Document Hub MUST honour any MFA gate enforced by Access Manager before allowing those transitions.

10. Integration Summary

- **Access Manager** — synchronous permission checks on every request; real-time revocation enforcement; RBAC rule definition
- **Digital Forms** — inbound signed-form PDFs at generation time, with full signing metadata (FormType, PatientId, SignatureTimestamp, DelegatedSigningAttribution, SignedPdfReference); immutability enforced on receipt
- **Aftercare Manager** — inbound care documents at generation time
- **Hygiene Subscriptions** — inbound signed subscription agreement PDFs at generation time; immutability enforced on receipt; secure reference returned to Hygiene Subscriptions for patient-record linkage
- **Care Plan Subscriptions** — inbound signed care plan contract PDFs at generation time; immutability enforced on receipt; secure reference returned for patient-record linkage
- **Referral Manager** — inbound referral supporting documents and AI-drafted referral letters; inbound outcome reports via email ingestion; outbound secure references for all ingested referral artefacts
- **Task Manager** — outbound task triggers for unassigned documents and document-linked actions
- **Communication Hub** — outbound notification triggers; secure time-limited references issued on demand; Communication Hub MUST NOT store raw document content
- **Lab Manager** — inbound artefact association requests; outbound secure `LabArteFact` references
- **PMS (e.g. Dentally)** — outbound secure metadata and Primoro links; PMS MUST NOT store raw files
- **Security & Privacy module** — AuditEvent records extend SecurityEvent shape, including View, Share, Revoke, and Delete events; immutable append
- **Smart Dashboards** — downstream analytics consumer (read-only engagement signals; Document Hub does not own the dashboard surface)
- **AI Guardian** *(future extension)* — unusual access-pattern detection; human-review only
- **Governance Reporting** *(future extension)* — compliance exports
- **Group Controls** *(future extension)* — multi-site document libraries

11. Explicit Non-Goals

- Acting as a generic file share — Document Hub is a governed, RBAC-controlled system, not a team drive.
- Sending documents as email attachments — all external delivery uses secure, time-limited links.
- Duplicating documents in external systems — all integrations consume secure references only.
- Exposing raw storage URLs — storage references are internal; no public or pre-signed URLs are surfaced directly.
- Bypassing RBAC or audit controls under any circumstance.
- Owning or ingesting Smart Treatment Proposal PDFs — owned by Smart Treatment Proposals.
- Providing clinical viewers for advanced imaging — identified as a future extension (Imaging module) and explicitly out of scope for MVP.

12. Versioning & Governance

This specification is owned by: the Document Hub module owner.

Changes to this spec require:

- Review by the MVP module owner.

- Impact analysis across all declared related modules (see /propose), with particular attention to Access Manager, Communication Hub, Digital Forms, Lab Manager, Hygiene Subscriptions, Care Plan Subscriptions, and Referral Manager given the tight integration contracts declared in §6.
- Version bump (patch for clarifications, minor for additive capability, major for breaking contract changes).

All future changes must preserve:

- The single-source-of-truth principle (one master copy; all integrations by secure reference).
- The secure-by-reference integration pattern.
- Complete, immutable auditability across all document events.
- Patient control over third-party sharing.
- The Digital Forms signed-PDF immutability requirement.
- The immutability requirement for all module-generated signed artefacts (Hygiene Subscriptions, Care Plan Subscriptions).
- The Smart Treatment Proposals explicit exclusion.

13. Build Contract (Engineering & QA)

13.1 Canonical Data Model

```
Document (
 DocumentId          UUID PRIMARY KEY,
 Category            VARCHAR NOT NULL,
 PatientId           UUID NULL,
 Source              ENUM('Staff','Patient','Module','External') NOT NULL,
 CurrentVersionId    UUID NOT NULL REFERENCES DocumentVersion(VersionId),
 LifecycleState      ENUM('Draft','Approved','Superseded','Archived',
                          'DeletedPendingPurge','Purged') NOT NULL,
 CreatedBy           UUID NOT NULL,
 CreatedByRole       VARCHAR NOT NULL,
 CreatedAt           TIMESTAMPTZ NOT NULL
)

DocumentVersion (
 VersionId           UUID PRIMARY KEY,
 DocumentId          UUID NOT NULL REFERENCES Document(DocumentId),
 CreatedBy           UUID NOT NULL,
 CreatedTimestamp    TIMESTAMPTZ NOT NULL,
 FileHash            VARCHAR NOT NULL,
 StorageReference    VARCHAR NOT NULL
)

Annotation (
 AnnotationId        UUID PRIMARY KEY,
 DocumentVersionId   UUID NOT NULL REFERENCES DocumentVersion(VersionId),
 CreatedBy           UUID NOT NULL,
 Content             TEXT NOT NULL,
 Timestamp           TIMESTAMPTZ NOT NULL
)

ShareLink (
 LinkId              UUID PRIMARY KEY,
```

```
 DocumentId          UUID NOT NULL REFERENCES Document(DocumentId),
 CreatedBy           UUID NOT NULL,
 CreatedByType       ENUM('Staff','Patient') NOT NULL,
 ExpiryTimestamp     TIMESTAMPTZ NOT NULL,
 Revoked             BOOLEAN NOT NULL DEFAULT FALSE
)

AuditEvent (
 EventId             UUID PRIMARY KEY,
 EventType           ENUM('Upload','Ingest','View','Download','Annotate',
                          'Share','Revoke','Acknowledge','VersionChange',
                          'Delete','Purge','IntegrationNotification',
                          'Escalation') NOT NULL,
 ActorUserId         UUID NOT NULL,
 ActorRole           VARCHAR NOT NULL,
 ActorSessionId      UUID NULL,
 TargetDocumentId    UUID NULL,
 TargetVersionId     UUID NULL,
 TargetShareLinkId   UUID NULL,
 DeviceId            VARCHAR NULL,
 Timestamp           TIMESTAMPTZ NOT NULL
 -- immutable once written; no UPDATE or DELETE permitted
)
```

13.2 Core Behaviour Rules

1. Every document MUST have exactly one master copy in Document Hub storage; no raw file may be written to any external system.
2. Every access to a document or secure reference MUST be permission-checked against Access Manager in real time; permission decisions MUST NOT be cached beyond the scope of a single request.
3. Annotations MUST be stored as overlays and MUST NOT alter the original DocumentVersion file or its FileHash.
4. Superseded DocumentVersions MUST be retained and MUST remain immutable; they may not be deleted unless the parent Document reaches the Purged state after its retention period has elapsed.
5. Signed-form PDFs ingested from Digital Forms MUST be version-locked on ingestion; no subsequent operation may alter the stored file or its FileHash.
6. Signed subscription agreement PDFs ingested from Hygiene Subscriptions and signed care plan contract PDFs ingested from Care Plan Subscriptions MUST be treated as immutable upon ingestion; no subsequent operation may alter the stored file or its FileHash.
7. Smart Treatment Proposal PDFs MUST NOT be ingested; any integration path that would result in their ingestion MUST be blocked at the API layer.
8. All ShareLinks MUST carry an explicit ExpiryTimestamp; links with no expiry MUST be rejected at creation time.
9. Deletion of a document MUST synchronously revoke all associated ShareLinks (set Revoked = TRUE) and terminate all active viewing sessions before returning success to the caller.
10. When Access Manager signals a permission revocation or modification, Document Hub MUST terminate active viewing sessions for affected documents immediately and MUST return access-denied on any subsequent secure-reference resolution for that actor.
11. OCR processing MUST be asynchronous and MUST NOT block document ingestion or state promotion.

12. Lab artefacts from authorised scanners MUST be categorised as `LabArtefact` and MUST be exposed to Lab Manager via a secure reference; the raw file MUST NOT be transmitted to Lab Manager.

13. Communication Hub MUST receive only secure time-limited references; Document Hub MUST NOT transmit raw file content to Communication Hub.

14. Every state transition and every document action listed in §8 MUST produce an immutable AuditEvent record; the absence of an expected audit record is a build defect.

15. Digital Forms inbound payloads MUST include FormType, PatientId, SignatureTimestamp, DelegatedSigningAttribution (where applicable), and SignedPdfReference; ingestion MUST be rejected if mandatory fields are absent.

16. Staff App Mode acknowledgement confirmations MUST be recorded as immutable `Acknowledge` AuditEvents by Document Hub, capturing UserId, Role, SessionId, document version, site, and timestamp.

17. AuditEvent records for View, Share, Revoke, and Delete events MUST be emitted to the Security & Privacy module's platform-wide audit trail in the canonical SecurityEvent shape; failure to emit is a build defect.

13.3 Configuration Surfaces

Practice-level settings (Admin Control Plane):

- Document categories and metadata schema per category.
- Role access permissions per document category.
- Acknowledgement requirements per category or specific document.
- ShareLink default expiry durations.
- Retention policies per category (duration; soft-delete-to-purge schedule).
- OCR and full-text search behaviour (enabled / disabled per category).

Per-user preferences (Access Manager):

- Notification preferences for document delivery events are governed by Communication Hub and Access Manager; Document Hub defers to those modules.

Per-document overrides (Document Hub):

- Individual ShareLink expiry overrides at creation time (within the bounds of practice-level policy).
- Per-document acknowledgement requirement override, where permitted by the practice-level category configuration.

13.4 Filtering & Views

Standard user filters:

- Category
- Patient
- Date range
- Source (`Staff | Patient | Module | External`)
- Lifecycle state / version state

Manager-level filters (additional):

- Acknowledgement status (outstanding / completed) by user, role, or site
- Access events (by actor, document, or date range)
- Active external shares and revocation history

Search:

- Full-text search across OCR-indexed documents, respecting RBAC.
- Search-within-document highlighting in the embedded viewer.

13.5 Module Extension Map

Future extensions that MUST NOT break this build contract:

- **AI Guardian** — unusual document access-pattern detection; activates anomaly audit events; human-review only; no AI-autonomous action.
- **Governance Reporting** — compliance export surface; reads AuditEvent records; no write access to Document Hub objects.
- **Group Controls** — multi-site document library sharing; extends Category and RBAC configuration surfaces; MUST honour single-source-of-truth principle.
- **Imaging** — advanced clinical viewer for imaging artefacts; replaces embedded viewer for applicable categories; Document Hub retains storage and lifecycle ownership.

13.6 Acceptance Criteria

The build of Document Hub is complete when:

- [ ] Documents ingest from all defined sources (staff upload, patient submission, module-generated, scanner/email) and a single master copy is stored, encrypted, and indexed.
- [ ] Duplicate file storage in external systems is structurally prevented — no raw file paths or content are transmitted outside Document Hub.
- [ ] Full-text search works inside OCR-indexed documents and respects RBAC.
- [ ] Secure sharing is implemented with mandatory expiry, is auditable, and is revocable in real time.
- [ ] Acknowledgement tracking is functional per user, role, and site; document updates reset acknowledgements.
- [ ] All state machine transitions enforce the rules in §3.2, including immutability of signed Digital Forms PDFs.
- [ ] All canonical objects can be created, read, and updated through the API; all state transitions are recorded in AuditEvent.
- [ ] Active viewing sessions terminate immediately on permission revocation; access-denied is returned on subsequent reference resolution.
- [ ] Lab Manager artefacts are categorised as `LabArtefact`, linked to the correct Lab Case, and exposed only via secure reference.
- [ ] Communication Hub obtains document references via the secure-by-reference pattern; no raw files are stored outside Document Hub.
- [ ] Smart Treatment Proposal PDF ingestion is blocked at the API layer (negative test passes).
- [ ] Digital Forms inbound payloads are validated for required fields (FormType, PatientId, SignatureTimestamp, SignedPdfReference); incomplete payloads are rejected.

- [ ] Hygiene Subscriptions and Care Plan Subscriptions signed agreement PDFs are ingested, treated as immutable on receipt, and a secure reference is returned to the originating module (positive test passes).
- [ ] Referral Manager supporting documents, AI-drafted letters, and outcome reports are ingested and linked to referral records via secure reference; unidentifiable artefacts enter the unassigned queue.
- [ ] Staff App Mode acknowledgement confirmations are recorded as immutable `Acknowledge` AuditEvents with the full actor and document-version context.
- [ ] AuditEvent records for View, Share, Revoke, and Delete events are emitted to the Security & Privacy module's platform-wide audit trail in the canonical SecurityEvent shape (positive and negative tests pass).
- [ ] AI boundaries in §7 are enforced (negative tests pass).
- [ ] The complete audit trail defined in §8 is captured for every event type; no event type is missing.
- [ ] Access control is enforced per §9, including real-time permission checks and MFA gate passthrough where Access Manager requires it.
- [ ] All non-functional requirements in §14 are met.

14. Non-Functional Requirements

- **Performance:** Document ingestion (upload receipt and encryption) MUST complete within an acceptable interactive latency. OCR processing MUST be asynchronous and MUST NOT block ingestion or viewer access. The embedded viewer MUST stream documents responsively without writing local copies.
- **Reliability:** No document loss under any partial failure condition. Storage writes MUST be durable before the ingestion success response is returned. Safe recovery from partial failures MUST be documented and tested. Availability target to be defined by the platform engineering team prior to production release.
- **Scalability:** Document Hub MUST support multi-site and multi-tenant operation, with category-level RBAC and retention policies independently configurable per practice and site. The data model and storage layer MUST be validated for expected document volume at group-practice scale before production launch.
- **Security:** All documents MUST be encrypted at rest and in transit. No public or raw storage URLs may be exposed. All external access requires Primoro authentication. Zero-trust: permission is checked on every request, never assumed from session state. Key management and secrets handling MUST follow platform security standards. Penetration testing of the secure-reference issuance and revocation flow is required before production release.
- **Privacy:** Document Hub MUST honour GDPR rights including right of access (exportable audit log), right to erasure (soft-delete → retention → purge schedule), and data minimisation (no unnecessary duplication). Retention policies MUST be configurable to meet CQC and healthcare governance requirements. All deletion and purge actions are logged immutably.
- **Observability:** Document Hub MUST export the following to the platform observability stack: ingestion throughput and error rates; OCR processing queue depth and latency; viewer session counts; ShareLink creation and access counts; permission-check latency; AuditEvent write latency and volume. Distributed tracing MUST be instrumented across all cross-module calls (Access Manager checks, outbound reference issuance, AuditEvent writes).
- **Accessibility:** The embedded viewer and all Document Hub UI surfaces MUST meet WCAG 2.1 AA accessibility standards, consistent with platform-wide accessibility requirements.

15. Open Questions

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

1. **Retention policy defaults** — The original states that retention policies are configurable but does not specify default retention durations for any document category. What are the minimum default retention periods, particularly for signed clinical documents, to satisfy CQC and GDPR obligations?

2. **Asynchronous OCR completion signalling** — The spec requires OCR to be asynchronous, but does not define how the UI communicates OCR-in-progress state to users, or what search behaviour is expected for documents whose OCR has not yet completed. How should partially-indexed documents be treated in search results?

3. **Manual patient-assignment queue — ownership and escalation** — Unassigned scanner and email artefacts are held in a queue for manual review. It is not specified which role owns that queue, how items are surfaced (e.g. via Task Manager task generation or a Document Hub-native view), or what the escalation path is if items remain unassigned beyond a threshold. These decisions must be made before the ingestion flow can be fully built.

4. **Availability target** — No specific availability SLA is stated (e.g. 99.9% uptime). This must be defined before the platform engineering team can design the storage and recovery architecture.

5. **Group Controls and multi-site document libraries** — The extension map identifies Group Controls as a future capability, but the scope of cross-site document sharing (which categories may be shared, who may initiate sharing, how RBAC interacts with the single-source-of-truth principle) is not defined. This boundary must be clarified before Group Controls integration is scoped.

6. **MFA gate specificity** — The spec delegates MFA requirements for sensitive operations to Access Manager but does not declare which Document Hub operations (delete, purge, bulk-share, or others) are expected to require MFA gates. The list of MFA-gated operations should be agreed with the Access Manager team and declared explicitly in this spec.

7. **Referral document category taxonomy** — The spec introduces a referral-document category for Referral Manager artefacts but does not define sub-categories (e.g. distinguishing supporting clinical documents from AI-drafted referral letters from inbound outcome reports). The taxonomy should be agreed with the Referral Manager module owner before build begins.