SANITIZED PORTFOLIO EXTRACT

Enterprise Knowledge & Onboarding Platform

Solution Architecture Brief · Governed retrieval and workflow guidance for regulated pharmaceutical content operations

DOCUMENT ID	SA-KEP-001	**VERSION / STATUS**	0.9 · Architecture review
PREPARED BY	Christopher Mangun · Solution Architecture / Forward-Deployed Delivery	**REVIEW AUDIENCE**	Platform Engineering, Security, MLR Operations, Product Ownership
SYSTEM OWNER	Enterprise Content Operations	**REVIEW DATE**	June 5, 2026 · Reconstructed from sanitized engagement artifacts

DECISION REQUESTED

Approve the proposed target architecture and pilot controls for a Teams-based knowledge assistant that retrieves approved enterprise content, provides cited process guidance, and preserves existing MLR approval authority.

1. PURPOSE AND CONTEXT

A global pharmaceutical enterprise needed to onboard a large delivery organization during a compressed operating transition. Content producers and reviewers relied on more than 20 disconnected repositories, collaboration spaces, and workflow systems. Staff frequently selected outdated claims or incorrect templates, increasing search effort, rework, and first-review rejection.

The proposed platform, referred to here as the Enterprise Knowledge Platform, provides one governed access layer over existing systems of record. It does not replace Veeva, SharePoint, work-management platforms, or MLR review. It improves discovery, preserves source authority, and guides users through controlled processes.

Engagement-derived conditions

- Approximately 1,000 users required accelerated onboarding.
- More than 20 content and workflow sources were in use.
- Search commonly required manual navigation across repositories.
- MLR review remained the authoritative approval mechanism.

Architecture principles

- Systems of record remain authoritative.
- Every answer identifies its approved source.
- Model output is treated as untrusted until validated.
- Uncertain or non-compliant outputs are blocked or escalated.

2. SCOPE, ASSUMPTIONS, AND EXCLUSIONS

In scope	Teams interaction, enterprise identity, tenant-aware retrieval, approved-content indexing, citations, process guidance, policy validation, audit logging, operational telemetry, and controlled pilot rollout.
Out of scope	Autonomous MLR approval, source-content authoring, replacement of source repositories, direct publication of regulated content, clinical decision support, and unrestricted web retrieval.
Assumptions	Source-system permissions are reliable; approved-content status is available as metadata; enterprise identity claims can be mapped to brand and regional access; source owners remain accountable for content quality.
Dependencies	Microsoft Teams and Entra ID, approved repository APIs, Veeva and SharePoint access, search/indexing services, enterprise logging, secrets management, and named operational support ownership.

3. SYSTEM CONTEXT

Content producer / new hire
Searches approved material and requests process guidance through Teams.

Enterprise Knowledge Platform
Authenticates, retrieves, validates, cites, logs, and returns governed responses.

MLR reviewer / content owner
Maintains approval authority and resolves exceptions or uncertain guidance.

IDENTITY AND QUERY ↓ · CITED ANSWER AND ESCALATION STATUS ↑

Systems of record
Veeva Vault, SharePoint, work management, CLM, regional repositories.

Enterprise controls
Identity, private networking, secrets, audit logging, retention, SIEM.

Operational services
Search index, model endpoint, telemetry, incident and support workflows.

Trust changes at four boundaries: user identity, source authorization, model inference, and downstream workflow action.

4. PROPOSED ARCHITECTURE

EXPERIENCE	**Microsoft Teams interface** Conversational access and guided onboarding	**Enterprise identity** SSO, role, brand, and regional claims	**Human escalation** MLR and content-owner review queues
APPLICATION AND CONTROL	**Request orchestrator** Intent, workflow state, timeout, and retry handling	**Validation gateway** Schema, citation, confidence, and policy checks	**Audit service** Query, sources, model version, result, and disposition
KNOWLEDGE AND AI	**Hybrid retrieval** Keyword, vector, metadata, and tenant filtering	**Relationship index** Brand, market, claim, template, and process relationships	**Model endpoint** Grounded synthesis only; no source authority
DATA AND PLATFORM	**Connector services** Incremental ingestion and approval metadata	**Private platform services** AKS, search, storage, key management	**Operations** Metrics, traces, SIEM, backup, and recovery

IDENTITY BOUNDARY	CONTENT BOUNDARY	MODEL BOUNDARY	APPROVAL BOUNDARY

5. PRIMARY REQUEST FLOW

STEP	COMPONENT	BEHAVIOR	FAILURE / CONTROL RESPONSE
1	Teams + identity	Authenticate the user and attach brand, market, and role claims.	Reject unauthenticated requests; do not create an anonymous retrieval path.
2	Request orchestrator	Classify the request as search, process guidance, or unsupported action.	Route unsupported or consequential actions to a human-owned workflow.
3	Retrieval service	Filter by tenant and approval metadata before keyword/vector retrieval.	Return no answer when authorized approved evidence is insufficient.
4	Model endpoint	Synthesize an answer from bounded retrieved context.	Model output remains untrusted and cannot directly update systems of record.
5	Validation gateway	Verify citations, response schema, policy rules, and confidence threshold.	Block, request clarification, or escalate failed validation.
6	Audit service	Record identity scope, sources, versions, checks, response, and disposition.	Fail closed for workflows requiring evidence capture.

6. ARCHITECTURE DECISIONS

ID	STATUS	DECISION	RATIONALE AND CONSEQUENCE
ADR-01	**ACCEPTED**	Keep source repositories authoritative; build an access layer rather than migrate content into a new authoring system.	Reduces change-management and validation scope. Requires reliable connector health and metadata quality.
ADR-02	**ACCEPTED**	Use hybrid retrieval with tenant, brand, market, and approval-state filters applied before model context construction.	Improves precision and access control. Adds indexing complexity and requires explicit metadata ownership.
ADR-03	**ACCEPTED**	Require citations for every substantive answer and block responses without approved evidence.	Improves reviewability and user trust. Some queries return no answer even when a plausible response could be generated.
ADR-04	**ACCEPTED**	Place validation and policy checks outside the prompt and model endpoint.	Creates deterministic enforcement and testability at the cost of additional latency and operational components.
ADR-05	**PILOT**	Use Teams as the initial experience channel.	Reduces adoption friction and reuses enterprise identity. Couples the initial experience to Microsoft collaboration services.
ADR-06	**DEFERRED**	Do not permit autonomous MLR submission or approval during the initial release.	Limits early automation value but keeps regulatory accountability and workflow authority unchanged during pilot validation.

7. NONFUNCTIONAL REQUIREMENTS

ID	QUALITY	TARGET	ARCHITECTURE RESPONSE	VERIFICATION
NFR-01	Availability	99.9% monthly service target	Multi-zone application services, health probes, graceful “search unavailable” response, and connector isolation.	Availability dashboard, synthetic query, quarterly recovery exercise.
NFR-02	Performance	Retrieval p95 <1s; answer p95 <5s	Bounded retrieval, index partitioning, caching of stable guidance, streaming response where appropriate.	Representative load test and production percentile traces.
NFR-03	Access control	No cross-brand or cross-market retrieval	Identity claims mapped to index filters; source permissions retained in indexed metadata.	Negative authorization tests and sampled access review.
NFR-04	Traceability	Reconstruct each user-visible answer	Persist query, source IDs, approval state, prompt/model version, validation results, and final disposition.	Evidence-reconstruction test for sampled responses.
NFR-05	Recovery	RTO <4h; RPO <1h	Geo-replicated configuration and indexes, backed-up audit data, documented failover and rebuild procedure.	Scheduled restore and regional failover exercise.

NFR-06	Safety	No uncited regulated guidance	Minimum evidence threshold, citation validation, unsupported-intent refusal, and human escalation.	Regression suite using approved, stale, conflicting, and unauthorized sources.

8. SECURITY, PRIVACY, AND COMPLIANCE CONTROLS

Preventive controls

- Enterprise SSO, MFA, role and market authorization.
- Private endpoints and restricted service identities.
- Encryption in transit and at rest with managed keys.
- Input validation and content-boundary filtering.
- Approved-source and approval-state metadata requirements.

Detective and corrective controls

- Central audit logging and SIEM forwarding.
- Retrieval-denial, citation-failure, and escalation metrics.
- Model and prompt version traceability.
- Connector-health alerts and stale-index detection.
- Documented disable, rollback, and evidence-preservation procedures.

The platform supports controlled content operations; it does not itself confer HIPAA, FDA, GxP, SOC 2, or 21 CFR Part 11 compliance. Applicable controls remain subject to enterprise validation and quality-system ownership.

9. DEPLOYMENT AND OPERATING MODEL

STAGE	ENTRY CRITERIA	EXIT CRITERIA	PRIMARY OWNER
Development	Synthetic and approved non-production content; isolated identities.	Unit, policy, retrieval, and evidence tests passing.	Platform Engineering
Controlled pilot	Security review complete; named content owners; support runbook approved.	Acceptance targets met; no material access-control or citation defects.	Product Owner + MLR Operations
Production rollout	Pilot sign-off; monitoring and incident ownership active; source SLAs agreed.	Phased user onboarding completed with stable SLOs and support demand.	Enterprise Content Operations

Operational telemetry

- Availability and end-to-end response latency
- Retrieval precision proxy and no-answer rate
- Citation validation and policy failure rates
- Escalation volume and time to resolution
- Index freshness and connector failures
- Model usage and cost by business area

Support and ownership

- L1: user access and known workflow guidance
- L2: content, indexing, and source-permission issues
- L3: orchestration, model, policy, and platform defects
- Content owners approve source corrections
- Security owns access and incident escalation
- Product owner approves scope and release changes

10. CAPACITY AND COST ASSUMPTIONS

ASSUMPTION	ARCHITECTURE PLANNING VALUE	VALIDATION REQUIRED BEFORE PRODUCTION
Users	Approximately 1,000 onboarded users; phased concurrency during rollout.	Measure peak active sessions and brand/region distribution during pilot.
Knowledge sources	20+ repositories and workflow systems with mixed update frequency.	Confirm API limits, metadata completeness, ownership, and refresh SLA per source.
Query profile	Search and process-guidance questions dominate; write actions excluded.	Classify pilot intents and validate unsupported-action rate.
Economics	Value case driven by reduced search, avoidable rework, and faster onboarding.	Use observed pilot volumes and loaded labor rates; do not treat portfolio estimates as approved business case.

11. RISKS AND OPEN DECISIONS

ID	RISK / DECISION	MITIGATION OR REQUIRED ACTION	OWNER	STATUS
R-01	Source metadata does not reliably identify approval state.	Define mandatory metadata contract and quarantine sources that do not meet it.	Content Governance	**OPEN**
R-02	Permissions differ between source systems and indexed content.	Reconcile identities and run cross-brand negative tests before pilot expansion.	Security / IAM	**OPEN**
R-03	Users may interpret cited guidance as formal approval.	Use clear response language, preserve MLR workflow, and route approval questions to reviewers.	MLR Operations	**MITIGATING**
D-01	Long-term model provider and hosting topology.	Complete security, latency, portability, and cost comparison after pilot evidence is available.	Architecture Board	**DEFERRED**
D-02	Retention period for prompts and generated responses.	Align business need, investigation requirements, privacy, and records policy.	Legal / Records	**REQUIRED**

12. REVIEW RECOMMENDATION

Proceed with a controlled pilot provided that source approval metadata, identity-to-content authorization, audit retention, and named support ownership are approved before user onboarding. Autonomous approval and direct write-back to regulated systems should remain outside the initial release.

Platform Engineering · Review / date

Security Architecture · Review / date

MLR Operations · Review / date

Provenance: reconstructed from a sanitized pharmaceutical enterprise knowledge-platform engagement. Organization names, implementation details, and selected measurements are generalized. Architecture structure and control decisions are engagement-derived; performance and economic figures remain planning targets unless explicitly validated in the case study.