

Integrating AI Governance into a Multi-Site Telepharmacy Platform

Architecture, Healthcare-Specific Design Decisions, and Lessons

Matthew J. Maughan, PharmD, MHCDS, CPEL Director of Telepharmacy, Dartmouth Health

The Problem

Telepharmacy platforms handle protected health information across dozens of hospital sites. AI agents are increasingly involved — verifying medications, routing orders, flagging interactions, generating reports. Every one of those actions touches PHI. Every one needs an audit trail.

The standard approach: CloudWatch logs with 7-year retention. It works for HIPAA compliance. But it doesn't answer the harder questions:

- Which actions were AI-generated vs. human-directed?
- Can you prove an AI agent followed governance rules on a specific patient interaction?
- If an agent starts behaving differently, would you know?
- If something goes wrong, can you reconstruct exactly what happened — cryptographically?

Charter governance was integrated into PharmLink — a telepharmacy platform designed for 34 hospital sites across 5 states — to answer these questions.

The Platform

PharmLink is a Flask/Next.js telepharmacy workflow management system:

- **Backend:** Python 3.11, Flask 2.3.3, PostgreSQL 15, Redis, Celery
- **Frontend:** React 18, Next.js 16, TypeScript, Tailwind CSS
- **Infrastructure:** AWS (ECS, RDS Multi-AZ, S3, CloudWatch), Docker, Terraform
- **Designed for:** 34 hospital telepharmacy sites across 5 states
- **Security:** JWT authentication, role-based authorization (PHARMACIST, MANAGER, CUSTOMER, ADMIN), MFA, CSRF protection, HIPAA-compliant audit logging

The platform already had a comprehensive audit trail — a 908-line middleware (`audit_logging.py`) that captures every PHI access event with before/after snapshots, PHI field masking, and CloudWatch integration.

What it didn't have: governance.

What We Built

Charter Governance Service

A central service (`governance.py`) that loads the project's `charter.yaml` at Flask startup and exposes three functions:

- `check_layer_a(action, context)` — validates against hard constraints (returns permit/deny)

- `check_layer_b(action, context)` — checks escalation thresholds (returns permit/escalate)
- `log_decision(action, rule, outcome, actor)` — writes to Charter’s immutable hash chain

The governance config defines the rules in plain English:

```
domain: healthcare_operations
governance:
  layer_a:
    rules:
      - Never send patient communications without pharmacist review
      - Never access patient records without explicit authorization
      - Never disclose PHI outside authorized channels
      - Never bypass MFA or authentication controls
      - Never deploy to production without staging validation
  layer_b:
    rules:
      - patient_communication: always requires human_approval
      - clinical_decision: always requires human_approval
      - data_access (sensitive): requires human_review
      - production_deployment: always requires human_approval
  kill_triggers:
    - hipaa_violation_detected
    - audit_trail_tampering
    - unauthorized_production_access
    - credential_exposure
```

Dual-Write Audit Trail

Every audit event goes to two destinations:

1. **CloudWatch** (existing) — HIPAA-compliant, 7-year retention, AWS BAA in place
2. **Charter hash chain** (new) — HMAC-SHA256 signed, immutable, cryptographically verifiable

CloudWatch remains the primary compliance trail. Charter chain is the secondary, cryptographically verifiable trail. This avoids making Charter a single point of failure for regulatory compliance.

The mapping: | Audit Action | Charter Event Type | |—————|—————| | VIEW |
 phi_access_view | | CREATE | phi_record_created | | UPDATE | phi_record_modified |
 | DELETE | phi_record_deleted |

Layer A Enforcement Middleware

A new middleware in the Flask stack, positioned after authentication but before route handlers:

Request → Error Handler → Audit Logging → Authentication → GOVERNANCE → Authorization → Route

Every request is checked against Layer A hard constraints. PHI endpoint access without authorization, patient communications without pharmacist role, MFA bypass attempts — all blocked before they reach the route handler.

Kill Trigger Detection

Four kill triggers monitored continuously, each with a graduated response calibrated to the threat level:

Trigger	Detection	Response	Severity
HIPAA violation	Unauthorized PHI access (403 on PHI endpoint)	Degrade to read-only, alert admins via WebSocket	Medium — could be misconfigured role, not attack
Audit tampering	UPDATE/DELETE on audit_log table	Hard kill of affected service, immediate alert, immutable chain record	Critical — active compromise of trust infrastructure
Unauthorized deploy	Deploy without staging flag	Block the deploy (running system unaffected)	Low — prevent bad deploy, don't disrupt care
Credential exposure	Pattern scan for API keys, JWT secrets in request/response	Hard kill of exposed surface, redact, alert, chain record	Critical — active credential compromise

Escalation Queue

Layer B decisions route to a human approval queue:

- Patient communications → pharmacist sign-off required
- Clinical decisions → licensed pharmacist approval
- Production deployments → admin authorization
- Database migrations → admin review

The queue includes WebSocket push for real-time notifications, 24-hour expiration, and full chain logging of approval/denial decisions.

Healthcare-Specific Design Decisions

These are the decisions that matter most — and the ones that would be different in a non-healthcare context.

Decision 1: Fail-Closed for Layer A, Fail-Open for Layer B

Layer	Behavior	Rationale
Layer A	Fail-closed	If governance can't verify, block the action. HIPAA violations are non-negotiable.

Layer	Behavior	Rationale
Layer B	Fail-open with logging	If governance service is unavailable, allow the action but log aggressively and alert. Blocking a pharmacist from verifying a medication during a patient interaction is worse than a missed escalation.

This is the single most important design decision. In healthcare, blocking a clinician from doing their job can harm patients. The governance system must never become a barrier to patient care.

Decision 2: Graduated Response, Not Binary Kill

Kill triggers in healthcare software cannot all produce the same response. A telepharmacy platform designed for 34 hospitals requires nuance — the response must match the threat.

- **Critical threats (audit tampering, credential exposure):** Hard kill of the affected service. If someone is modifying the audit trail or credentials are compromised, the system cannot be trusted. Shut down the exposed surface immediately.
- **Medium threats (unauthorized PHI access):** Degrade to read-only mode. Could be a misconfigured role, not an attack. Alert admins, log to chain, but don't block pharmacists from seeing orders.
- **Low threats (unauthorized deploy):** Block the specific action. The running system is unaffected — just prevent the bad deploy.

In all cases: alert every admin via WebSocket, log to the immutable chain, and the pharmacist can still see the order and verify the medication. Patient care is never interrupted by a governance response to a non-critical threat.

Decision 3: Chain as Secondary, Not Primary

CloudWatch remains the primary audit trail for HIPAA compliance. Charter chain is additive — it provides cryptographic verification, behavioral profiling, and pattern detection that CloudWatch cannot. But if Charter chain fails, the HIPAA audit trail is unaffected.

This decision was driven by a simple question: “If the governance system goes down at 2 AM and a pharmacist needs to verify a critical medication, what happens?” The answer must be: “Nothing changes for the pharmacist. The audit trail continues. We fix governance in the morning.”

Decision 4: Patient Safety Override (With Safeguards)

If a pharmacist explicitly invokes a patient safety override (e.g., emergency medication verification), Layer B escalations are bypassed. The override exists because blocking patient care is worse than the risk of misuse.

The risk is managed through multiple layers, not by eliminating the override:

- **MFA-verified credentials required.** PharmLink enforces multi-factor authentication on all pharmacist accounts. An override can only be triggered by a user who has passed MFA —

a stolen password alone is not sufficient. The attacker would need both the credentials and the authenticator device.

- **Immutable chain logging.** Every override is logged with `override_reason` (a required field) and a full chain record. The override is auditable, not invisible.
- **Immediate admin notification.** Override events trigger higher-visibility alerts than normal actions — admins are notified in real time via WebSocket.
- **Rate limiting.** More than 2-3 overrides per hour from the same user triggers a kill trigger investigation.
- **Mandatory post-hoc review.** Every override must be reviewed by a second pharmacist within 24 hours or it is automatically flagged for compliance investigation.

This is standard in clinical systems — the override mechanism is not a vulnerability when it is visible, rate-limited, MFA-gated, and subject to mandatory peer review.

HIPAA Control Mapping

Charter’s `charter.yaml` maps directly to HIPAA requirements:

HIPAA Requirement	Charter Implementation
Access Controls (164.312(a))	Layer A: explicit authorization required for PHI
Audit Controls (164.312(b))	Dual-write: CloudWatch + Charter hash chain
Integrity Controls (164.312(c))	HMAC-SHA256 signed chain entries, tamper detection
Transmission Security (164.312(e))	TLS enforced, PHI masking in logs
Minimum Necessary (164.514(d))	Layer B: sensitive data access requires human review
Breach Notification (164.400-414)	Kill trigger: credential exposure → immediate alert

Analytics: What the Chain Reveals

With Charter v3.2.0’s analytics engine, the governance chain becomes queryable:

- **Actor profiling:** Which users access which patient records, session patterns, temporal distribution
- **Anomaly detection:** Flags when access patterns deviate from baseline (e.g., a user suddenly accessing records outside their normal hospital assignment)
- **Sequence mining:** Discovers common workflow patterns (e.g., “view order → verify drug interaction → approve → document”) and flags deviations
- **Governance dashboards:** One-call summary of escalation rates, approval times, violation counts, chain health

All analytics run locally. No PHI leaves the machine. DuckDB provides millisecond queries over the chain data.

Results

Metric	Value
PHI access events dual-written to chain	100%
Layer A violations blocked at runtime	100% (fail-closed)
Layer B escalations routed to approval queue	100% for clinical/communication actions
Kill trigger detection to admin alert	< 5 seconds
Regression in existing functionality	Zero (all existing tests pass)
Additional infrastructure cost	\$0 (local-first, runs on existing servers)
Governance config file	64 lines of plain-English YAML

What We Learned

1. **Governance doesn't slow things down.** The middleware adds < 5ms per request. Pharmacists don't notice. The dual-write to chain is fire-and-forget (asynchronous). The system is no slower with governance than without.
2. **Plain-English rules matter.** The `charter.yaml` file is readable by pharmacists, compliance officers, and administrators — not just engineers. When a compliance officer asks “what are the rules?”, you hand them the YAML file. It reads like a policy document.
3. **The chain answers questions CloudWatch can't.** “Can you prove this AI agent followed governance rules on this specific patient interaction last Tuesday?” CloudWatch gives you logs. Charter gives you a cryptographically signed chain entry with the exact rule that was applied, the actor, the decision, and a hash linking it to every prior action.
4. **Healthcare governance requires nuance.** Fail-closed for safety, fail-open for care. Graduated kill responses matched to threat level. MFA-gated overrides with mandatory peer review. These aren't generic governance decisions — they're clinical workflow decisions that happen to be implemented in code.

Getting Started

```
pip install charter-governance
charter init --domain healthcare
charter generate
```

Three commands. 64 lines of YAML. Every AI agent in your healthcare platform gets boundaries, escalation paths, and a cryptographically verifiable audit trail.

Charter is open source (Apache 2.0): <https://github.com/germpharm/charter> **Website:** <https://charteragent.ai> **Contact:** matthew@charteragent.ai

Charter Governance v3.2.0 | 51 Python modules | 55 MCP tools | 8 compliance frameworks | Apache 2.0