

FHIR/TEFCA Clinical Workflow Decision Layer Playbook

Cover Page

Industry: Healthcare Document Type: Technical Solution Engineering White Paper Publisher: Kettle Logic Author: Matthew Loschiavo, Founder/CEO, Kettle Logic Editor: Matthew Loschiavo (Editorial Review) Version: v9.0 Published: 2026-02-24 Audience: CTO / VP Engineering / Solution Engineering / Platform / Security / Data

Document Control

Field	Value
Document ID	KL-HEALTHCARE-TECH-V9
Status	Published
Review Cadence	Quarterly or on major regulatory / technology change

Executive Summary

This paper focuses on Prior authorization triage and evidence assembly in Healthcare. The strategy keeps existing enterprise platforms as systems of record while building a governed system of decision for policy checks, scoring, AI assistance, and exception routing. The objective is measurable gains in revenue, cost, and risk reduction with stronger controls and lower future integration cost.

This v9 pass fixes repeated paragraphs and adds concrete artifacts: industry-specific KPI baseline/target ranges, pseudo-code policy rules, and a sample JSON event payload for a key workflow.

Table of Contents

1. Cover Page
2. Document Control
3. Executive Summary
4. Table of Contents
5. Business Decision Drivers
6. System Landscape Reality Check
7. System of Record vs System of Decision
8. Industry Workflow Focus
9. Industry-Specific KPI Baselines and Targets
10. Executive Strategy (5-Year / 10-Year)
11. Board/CFO Capital Allocation Lens
12. Technology Fit Matrix
13. Solution Architecture / Implementation Playbook
14. Sample Policy Rules (Pseudo-code)
15. Sample JSON Event Payload
16. AI Strategy and Governance
17. Privacy, GDPR, and Data Rights Constraints
18. Risk Register
19. Roadmap and Governance Cadence
20. Glossary
21. References
22. Appendices

Business Decision Drivers

Businesses make modernization decisions to increase revenue, reduce costs, and reduce risk. Programs also succeed or fail based on speed, resilience, and strategic optionality. A strong white paper translates technology decisions into

these business outcomes rather than relying on generic transformation language.

Primary motivations

- Revenue: throughput, conversion, coverage, retention, margin quality
- Cost: labor productivity, defect/rework reduction, dispute handling, runtime efficiency
- Risk: privacy, cyber, fraud, compliance, operational resilience, model risk

Additional motivations that often matter

- Time-to-market and change velocity
- Executive trust in controls and evidence
- Vendor portability and strategic flexibility

System Landscape Reality Check

ERP, CRM, POS, EHR, core admin, MES, SCADA, PIM/PXM/MDM, and WMS are not obsolete just because AI is new. In most organizations they remain the legal or operational source of truth. What changes is where high-speed decisions and policy enforcement should happen.

Reality-based strategy

- Preserve core stability and integrity
- Expose events/APIs and data quality telemetry
- Move decision logic into a governed layer
- Keep privacy and audit evidence attached to workflow decisions

System of Record vs System of Decision

SoR: authoritative transactions, master data, legal history
SoD: policy evaluation, AI recommendations, optimization, routing
SoX: operator queues, portals, partner APIs, copilots

Separating SoR from SoD reduces blast radius, improves reuse, and creates a practical path for staged capital allocation.

Industry Workflow Focus

Key workflow: Prior authorization triage and evidence assembly

In Healthcare, workflow modernization is often framed as a platform gap, but the real bottleneck is unclear thresholds. A stronger approach starts with one workflow, one KPI stack, and one policy owner so teams can prove value without destabilizing core systems.

The practical modernization challenge in Healthcare is not lack of software; it is inconsistent decisions around policy governance. When thresholds, routing rules, and exception ownership vary by team, cycle time and defect costs rise even if all major systems are present.

For Healthcare operators, decision automation becomes useful only when it changes execution behavior. That requires explicit policy traces, queue prioritization, and evidence packets that supervisors can review-not just a dashboard or a model score.

Leaders in Healthcare should evaluate exception routing as a control-and-economics problem. The win condition is not maximum automation; it is faster, safer decisions with measurable improvements in revenue, cost, and risk metrics.

A durable Healthcare strategy for AI-assisted triage avoids two traps: broad core replacement before ROI is proven, and AI-first pilots with weak governance. The recommended pattern is a governed decision layer with clear SoR boundaries, policy versioning, and staged autonomy.

In Healthcare, operating discipline is often framed as a platform gap, but the real bottleneck is missing KPI baselines. A stronger approach starts with one workflow, one KPI stack, and one policy owner so teams can prove value without destabilizing core systems.

The practical modernization challenge in Healthcare is not lack of software; it is inconsistent decisions around portfolio sequencing. When thresholds, routing rules, and exception ownership vary by team, cycle time and defect costs rise even if all major systems are present.

For Healthcare operators, evidence design becomes useful only when it changes execution behavior. That requires explicit policy traces, queue prioritization, and evidence packets that supervisors can review-not just a dashboard or a model score.

Leaders in Healthcare should evaluate queue management as a control-and-economics problem. The win condition is not maximum automation; it is faster, safer decisions with measurable improvements in revenue, cost, and risk metrics.

A durable Healthcare strategy for change control avoids two traps: broad core replacement before ROI is proven, and AI-first pilots with weak governance. The recommended pattern is a governed decision layer with clear SoR boundaries, policy versioning, and staged autonomy.

Industry-Specific KPI Baselines and Targets

These sample ranges are intended for planning and executive discussion. Final targets should be calibrated using your actual baseline, product/channel mix, and regulatory constraints.

KPI	Typical Baseline Range	Program Target Range	Business Driver
Prior auth turnaround	2-10 days	4-24 hours	Revenue / access
Denial rate	8-20%	4-12%	Revenue / risk
Documentation time per case	20-70 min	8-30 min	Cost / clinician burden
Care gap closure rate	25-55%	45-75%	Quality
Referral leakage	12-35%	5-18%	Revenue

KPI usage guidance

Use a balanced KPI set. Growth-only programs can quietly increase risk. Risk-only programs can become compliance-heavy and lose support. A monthly review should include at least one KPI from each column: growth, cost, and risk.

Executive Strategy (5-Year / 10-Year)

5-Year plan

Build reusable decision-platform capabilities (policy, workflow, observability, privacy, audit) and apply them to a small set of high-value workflows with visible KPI movement. Avoid broad multi-year replacement programs before workflow-level ROI is proven.

10-Year plan

Operate with stable systems of record and fast, governed systems of decision. Use a technology fit matrix to evaluate AI, blockchain, spatial/digital twin, and confidential computing based on workflow fit-not trend pressure.

Board/CFO Capital Allocation Lens

Treat modernization as a staged investment portfolio. Fund a 90-day proof phase, then a 12-month expansion phase, then platform reuse only when the economics and control evidence are visible.

Funding questions for executives

1. Which KPI improved and by how much?
2. Which costs were removed vs shifted?
3. What controls are now automated and testable?
4. What reusable assets (policies, contracts, events, runbooks) were created?

Technology Fit Matrix

Technology Pattern	Use Now / Pilot / Watch	Why	Typical Failure Mode
Data contracts + policy-as-code	Use now	Highest leverage for quality, controls, and reuse	Treated as docs, not enforced in tests
Bounded AI in workflows	Use now (gated)	Speeds triage and evidence assembly	No action classes / weak audit trail
Confidential computing	Pilot selectively	Good for regulated / sensitive collaboration	Added complexity without workflow fit
Spatial / digital twin	Pilot workflow-first	Strong for simulation and planning	Demo-driven instead of KPI-driven
Blockchain / shared ledger	Pilot selectively	Works for multi-party trust/provenance	Used where internal governance is the issue
PQC / crypto-agility	Plan now	Long-horizon risk reduction	Deferred until emergency migration

Solution Architecture / Implementation Playbook

Reference implementation sequence

1. Baseline KPI and map current exception types
2. Define SoR/SoD boundary for the selected workflow
3. Create a minimal event schema and data contract
4. Implement initial policy rules and evidence logging
5. Add bounded AI (assist/recommend) with approval gating
6. Publish operator runbooks and escalation paths
7. Instrument business + technical + cost telemetry

Architecture must-haves

- Correlation IDs across all workflow steps
- Policy and model versioning
- Idempotent event handling and replay safety
- Privacy tags and retention controls
- Explainable operator-facing decisions

Sample Policy Rules (Pseudo-code)

The sample below shows how business thresholds, privacy constraints, and exception routing can be encoded directly in the workflow control plane.

```
RULE PriorAuthTriage
WHEN request.procedure_code in high_risk_codes
THEN automation_class = "AssistOnly"
     AND require_clinician_review = true

WHEN PHI_scope_requested not_subset_of minimum_necessary_scope
THEN action = "DenyAccessAndLog"

WHEN payer.policy_match_confidence >= 0.90 AND evidence_bundle_complete == true
THEN route_queue = "FastTrackSubmission"
ELSE route_queue = "ManualReview"

WHEN patient.consent_status != "valid"
THEN block_external_exchange = true
```

Sample JSON Event Payload

This example payload illustrates the minimum structure needed for observability, auditability, and replay-safe workflow processing.

```
{
  "eventType": "PriorAuthCaseTriaged",
  "eventVersion": "1.0",
  "caseId": "PA-2026-009117",
  "patientId": "pt_7f3a2",
  "memberId": "m_99127",
  "payer": "PAYER-ALPHA",
  "procedureCode": "27447",
  "policyMatchConfidence": 0.93,
  "evidenceBundleComplete": true,
  "automationClass": "AssistOnly",
  "decisionQueue": "FastTrackSubmission",
  "phiScope": "prior_auth_required",
  "policyVersion": "hc.pa.v10",
  "modelVersion": "payer-policy-match-2.8",
  "evaluatedAt": "2026-02-24T17:04:10Z",
  "correlationId": "hc-e9fa22"
}
```

Event payload design notes

- Include eventVersion, policyVersion, and (if applicable) modelVersion
- Include entity IDs and correlationId
- Prefer references/tags over raw sensitive payloads when possible

- Ensure consumers can handle schema evolution safely

v10.1 Technical Interface Addendum

Sample API Endpoints and Request/Response Examples

Triage

POST /v1/healthcare/prior-auth/triage

Request

```
{
  "caseId": "PA-2026-009117",
  "payer": "PAYER-ALPHA",
  "procedureCode": "27447",
  "phiScopeRequested": "prior_auth_required",
  "evidenceBundleComplete": true
}
```

Response

```
{
  "decisionQueue": "FastTrackSubmission",
  "automationClass": "AssistOnly",
  "policyVersion": "hc.pa.v10"
}
```

Submit

POST /v1/healthcare/prior-auth/submit

Request

```
{
  "caseId": "PA-2026-009117",
  "submissionChannel": "payer_api",
  "attestationBy": "clinician_32"
}
```

Response

```
{
  "status": "Submitted",
  "payerReference": "ALPHA-772991",
  "correlationId": "hc-e9fa22"
}
```

SQL and Event Schema Examples

SQL table (example)

```
CREATE TABLE prior_auth_case_decision (
  case_id TEXT PRIMARY KEY,
  payer_id TEXT NOT NULL,
  procedure_code TEXT NOT NULL,
  policy_match_confidence NUMERIC(4,3),
  evidence_bundle_complete BOOLEAN NOT NULL,
  decision_queue TEXT NOT NULL,
  automation_class TEXT NOT NULL,
  policy_version TEXT NOT NULL,
  model_version TEXT,
  evaluated_at TIMESTAMPTZ NOT NULL,
  correlation_id TEXT NOT NULL,
  phi_scope TEXT NOT NULL
);
CREATE INDEX idx_prior_auth_payer_time ON prior_auth_case_decision(payer_id, evaluated_at DESC);
```

Event schema contract (example)

```
{
  "eventType": "PriorAuthCaseTriaged",
  "required": [
    "eventType",
    "eventVersion",
    "caseId",
    "payer",
    "procedureCode",
    "decisionQueue",
    "automationClass",
    "policyVersion",
    "evaluatedAt",
    "correlationId"
  ]
}
```

```

  ],
  "optional": [
    "modelVersion",
    "policyMatchConfidence",
    "evidenceBundleComplete",
    "phiScope"
  ]
}

```

RACI by Industry

Role	RACI	Responsibility
Revenue Cycle Director	A	Owns turnaround and denial KPIs
Prior Auth Team Lead	R	Executes queue triage and submissions
Clinical Reviewer	R	Approves high-risk cases and attestation
Privacy Officer	C	Approves PHI scope and retention controls
Interoperability/FHIR Team	C	Maintains payer/EHR integration reliability
CIO/CMIO	I	Reviews throughput and safety controls

Legend: R = Responsible, A = Accountable, C = Consulted, I = Informed

AI Strategy and Governance

AI should start in bounded roles: classify, summarize, prioritize, and prepare evidence. Higher-impact actions should remain approval-gated until policy coverage, monitoring, and operator trust are mature.

AI governance controls

- Action classes (read / recommend / draft / route / approve / execute)
- Confidence thresholds + abstain behavior
- Human review for high-impact decisions
- Drift monitoring + business outcome monitoring
- Fallback paths and incident runbooks

Privacy, GDPR, and Data Rights Constraints

Privacy is a system design requirement, not a legal appendix. The decision layer must enforce minimization, purpose limitation, retention, and rights handling across raw and derived data, including logs and evidence stores.

Required controls

- Role- and purpose-based access
- Retention/deletion policies for logs, caches, and derived artifacts
- Data subject / consumer rights workflows where applicable
- Cross-border processing awareness
- Reviewable evidence exports

Risk Register

Risk	Impact	Control pattern
PHI leakage	Can degrade revenue, cost, or trust outcomes	Policy thresholds + workflow routing + monitoring + review cadence
unsafe automation	Can degrade revenue, cost, or trust outcomes	Policy thresholds + workflow routing + monitoring + review cadence
consent mismatch	Can degrade revenue, cost, or trust outcomes	Policy thresholds + workflow routing + monitoring + review cadence
interop outages	Can degrade revenue, cost, or trust outcomes	Policy thresholds + workflow routing + monitoring + review cadence
denial backlogs	Can degrade revenue, cost, or trust outcomes	Policy thresholds + workflow routing + monitoring + review cadence

Roadmap and Governance Cadence

First 90 Days

- Establish baseline KPI ranges and workflow ownership
- Implement initial event contract and policy set
- Launch assist/recommend AI mode with evidence logging

- Publish runbooks and escalation matrix

12-Month Plan

- Expand to adjacent workflows using shared patterns
- Add drift/cost telemetry and quarterly fit-matrix reviews
- Standardize policy and contract testing in CI/CD

Governance cadence

- Weekly: queue health, defects, SLA misses, overrides
- Monthly: KPI and business-case review (growth/cost/risk)
- Quarterly: control maturity and technology fit refresh

Glossary

- System of Record (SoR): authoritative operational or legal system
- System of Decision (SoD): policy/AI/workflow layer for governed decisions
- Policy-as-Code: versioned executable business rules
- Data Contract: tested schema and semantics between producers/consumers
- Correlation ID: shared ID used to trace a workflow across systems
- Strategic optionality: reduced future cost of adopting new tools/channels

References

1. FHIR workflow orchestration
2. TEFCA exchange constraints
3. prior auth operating model
4. PHI minimum necessary
5. NIST AI RMF
6. NIST Privacy Framework
7. NIST CSF 2.0
8. GDPR legal framework
9. CISA Secure by Design

Appendices

Appendix A: Why this version is more concrete

This v9 pass includes realistic KPI ranges, domain-specific policy examples, and JSON event payloads so executive strategy and solution engineering can align on something implementable.

Appendix B: Adoption checklist

- Executive sponsor and workflow owner named
- KPI baseline/targets approved
- Policy owner and review cadence assigned
- Event contract tested
- Privacy controls validated
- Runbooks and fallbacks documented

In Healthcare, operator adoption is often framed as a platform gap, but the real bottleneck is supervisor trust gaps. A stronger approach starts with one workflow, one KPI stack, and one policy owner so teams can prove value without destabilizing core systems.

The practical modernization challenge in Healthcare is not lack of software; it is inconsistent decisions around policy drift. When thresholds, routing rules, and exception ownership vary by team, cycle time and defect costs rise even if all major

systems are present.

For Healthcare operators, queue design becomes useful only when it changes execution behavior. That requires explicit policy traces, queue prioritization, and evidence packets that supervisors can review-not just a dashboard or a model score.

Leaders in Healthcare should evaluate runtime economics as a control-and-economics problem. The win condition is not maximum automation; it is faster, safer decisions with measurable improvements in revenue, cost, and risk metrics.

A durable Healthcare strategy for vendor posture avoids two traps: broad core replacement before ROI is proven, and AI-first pilots with weak governance. The recommended pattern is a governed decision layer with clear SoR boundaries, policy versioning, and staged autonomy.

In Healthcare, incident response is often framed as a platform gap, but the real bottleneck is fallback readiness. A stronger approach starts with one workflow, one KPI stack, and one policy owner so teams can prove value without destabilizing core systems.

The practical modernization challenge in Healthcare is not lack of software; it is inconsistent decisions around audit evidence. When thresholds, routing rules, and exception ownership vary by team, cycle time and defect costs rise even if all major systems are present.

For Healthcare operators, portfolio prioritization becomes useful only when it changes execution behavior. That requires explicit policy traces, queue prioritization, and evidence packets that supervisors can review-not just a dashboard or a model score.

Leaders in Healthcare should evaluate change management as a control-and-economics problem. The win condition is not maximum automation; it is faster, safer decisions with measurable improvements in revenue, cost, and risk metrics.

A durable Healthcare strategy for measurement discipline avoids two traps: broad core replacement before ROI is proven, and AI-first pilots with weak governance. The recommended pattern is a governed decision layer with clear SoR boundaries, policy versioning, and staged autonomy.