

Insurance AI Architecture and Implementation Center

Executive handbook for production reference designs, decision boundaries, and control patterns.

REFERENCE DESIGN	PRIMARY USE
Predictive AI production	Scores, classes, probabilities, forecasts, and rankings
Enterprise RAG assistant	Current, approved, access-controlled knowledge
Fine tuning lifecycle	Stable repeated model behavior
Agentic insurance workflow	Multi-step coordination and action
Enterprise AI control plane	Shared policy, evidence, monitoring, and lifecycle governance


Executive Architecture Principles

Choose the simplest architecture that can meet the insurance objective. Govern the complete system, including data, models, prompts, retrieval, tools, identities, human decisions, downstream effects, vendors, telemetry, and recovery.

QUESTION	DESIGN RESPONSE
What must the system produce?	Separate prediction, content, current knowledge, learned behavior, and action.
What can it change?	Make system authority explicit. Add approvals and recovery as impact and irreversibility increase.
What evidence proves it works?	Define task metrics, counter-metrics, limitations, failure tests, and acceptance thresholds before release.
Who remains accountable?	Assign business, technical, risk, data, security, compliance, and operational responsibilities.
How does it fail safely?	Engineer limits, handoffs, pause, rollback, reconciliation, incident response, and retirement.


Predictive AI Production

Estimate a score, class, probability, forecast, or rank from governed data.



Enterprise RAG Assistant

Answer or draft using current, approved, access-controlled enterprise knowledge.




Insurance fit

Policy and procedure guidance, claim research, regulatory Q&A, underwriting manuals, service knowledge.

DECISION POINT	GUIDANCE
Choose when	Use when knowledge changes, must be cited, is proprietary, or differs by product, state, policy version, or role.
Human boundary	Users review consequential answers and citations; content owners approve sources; high-impact guidance routes to qualified reviewers.
Failure modes	Stale or poisoned sources, access leakage, wrong jurisdiction, poor retrieval, unsupported synthesis, or indirect prompt injection.
Build versus buy	Buy foundation models and search components when useful; retain source authority, access design, retrieval evaluation, and response policy.

Fine Tuning Lifecycle

Change stable model behavior using governed examples rather than storing current facts in model weights.




Insurance fit

Specialized classification, consistent extraction, controlled output formats, domain terminology and repeated task behavior.

DECISION POINT	GUIDANCE
Choose when	Use after prompting and retrieval are insufficient, the behavior is stable, and representative training and evaluation examples exist.
Human boundary	Domain experts curate examples, reviewers approve evaluation criteria, and release owners accept regressions and residual risk.
Failure modes	Memorization, hidden dataset defects, bias, regression, stale learned facts, weak generalization, or base-model change.
Build versus buy	Prefer managed training infrastructure unless control or economics require otherwise; retain dataset rights, evaluation, approval, and version lineage.

Agentic Insurance Workflow

Plan and coordinate multi-step work through tools, state, approvals, and bounded authority.




Insurance fit

Claim intake coordination, evidence gathering, task routing, customer follow-up, recovery and service workflows.

DECISION POINT	GUIDANCE
Choose when	Use only when action and coordination create enough value to justify stronger identity, authorization, observability, and recovery controls.
Human boundary	People approve high-impact or irreversible actions, resolve ambiguity, receive failed handoffs, and can pause or terminate a run.
Failure modes	Excess authority, compounding errors, loops, hidden side effects, credential misuse, cross-system inconsistency, or failed recovery.
Build versus buy	Buy orchestration components selectively; retain tool contracts, permissions, approval policy, run evidence, incident response, and operating ownership.

Enterprise AI Control Plane

Apply shared governance, security, evidence, evaluation, and operational controls across AI architectures.



Insurance fit

A common layer for claims, underwriting, fraud, servicing, distribution, pricing, and enterprise operations.

DECISION POINT	GUIDANCE
Choose when	Use as the shared foundation while preserving architecture-specific controls for models, retrieval, prompts, tools, and agents.
Human boundary	Business, risk, compliance, security, data, legal, and technology owners approve within defined responsibilities and escalation paths.
Failure modes	Inventory gaps, inconsistent controls, fragmented logs, unclear ownership, unapproved changes, or controls that exist only on paper.
Build versus buy	Standardize shared services where possible; keep accountable ownership, policy decisions, risk acceptance, and evidence requirements internal.

Insurance Workflow Patterns

WORKFLOW	ARCHITECTURE	HUMAN AUTHORITY
Claims intake and triage	Extraction plus predictive triage plus RAG	Human owns coverage and material disposition
Evidence coordination	RAG plus bounded tool-enabled agent	Approval before consequential external action
Underwriting workbench	Predictive score plus RAG plus drafting	Underwriter owns selection and exceptions
Fraud support	Predictive propensity plus analytics plus summary	Investigator owns referral and adverse action
Policy servicing	RAG assistant or approved transaction agent	Read-only default; approve material changes

Release Readiness

GATE	REQUIRED EVIDENCE
Purpose and ownership	Approved use, users, affected decisions, owners, jurisdictions, prohibited uses
Architecture	Data flow, components, versions, identities, permissions, tools, vendors, downstream systems
Validation	Task metrics, robustness, fairness, security, privacy, human factors, failure and recovery
Operations	Monitoring, support, incident response, rollback, change control, reassessment, retirement

Implementation Roadmap

- 1 Frame the insurance outcome and authority.
- 2 Choose the smallest sufficient architecture.
- 3 Map data, models, prompts, retrieval, tools, identities, vendors, and downstream effects.
- 4 Set the risk tier, evidence standard, and owners.
- 5 Build evaluation, oversight, security, telemetry, and recovery with the product.
- 6 Pilot with bounded users, data, actions, jurisdictions, and duration.
- 7 Release through documented acceptance and risk approval.
- 8 Operate, reassess material changes, and retire cleanly.

Sources

NIST AI Risk Management Framework

NIST Generative AI Profile

NAIC Artificial Intelligence resource center

Google Cloud MLOps architecture guidance

OpenAI practical guide to building agents