

INSURANCE

Agent Governance for Indian Insurance

A strategic guide to governing autonomous AI agents across underwriting, claims, and fraud detection. From IRDAI compliance to production agent operations.

EXECUTIVE SUMMARY

India's insurance penetration stands at 3.7% of GDP — roughly half the global average. Over 650 million Indians lack any form of life or health insurance. Insurers are deploying AI agents that autonomously underwrite policies, settle claims worth Rs 59,000 crore annually, and detect fraud in the estimated 30% of health claims that involve some element of it. Yet when an agent denies your health claim, nobody can explain why. IRDAI's AI/ML guidelines demand agent governance that most insurers haven't built. This guide provides the roadmap for closing that gap.

01 The Coverage Gap

Agents settling claims for 650M uninsured

02 Agents Without Operations

What ungoverned agents cause

03 Agent Use Cases

Underwriting, claims, fraud agents

04 Agent Economics

Cost architecture for claims agents

05 IRDAI Compliance

Regulatory framework for agents

06 Agent Operations Stack

Operations architecture for insurance

07 Production Readiness

Five gates for insurance agents

08 The Platform

Agent governance infrastructure

09 Agent Implementations

What insurers build

10 Getting Started

IRDAI Agent Governance Accelerator

Agents Are Settling Claims. Nobody's Explaining.

India has one of the lowest insurance penetration rates in the world. Insurers are responding by deploying AI agents that autonomously underwrite policies, assess claims, and detect fraud. But when an agent denies your claim, who explains why?

3.7%

INSURANCE PENETRATION - IRDAI ANNUAL REPORT, FY2023-24

650M+

UNINSURED INDIANS - SWISS RE INSTITUTE, 2024

30%

FRAUD IN HEALTH CLAIMS - FICCI-EY, 2023

Four Gaps That Define the Problem

The Penetration Gap

At 3.7% of GDP, India's insurance penetration is roughly half the global average. AI agents are being deployed to close this gap, but they're making decisions for populations that have never held a policy. No claims history. No actuarial baseline. No precedent for what fair looks like.

The Fraud Agent Gap

An estimated 30% of health insurance claims involve some element of fraud — from inflated bills to phantom surgeries. Hospital-insurer collusion rings bill for procedures never performed. Staged accident networks file identical motor claims across insurers. A single fraud model can't reason across this complexity.

The Claims Agent Gap

Indian insurers settled over Rs 59,000 crore in claims in FY24. Motor claims agents now auto-assess vehicle damage from photographs without human review. Health claims agents validate procedures against policy terms and deny coverage in milliseconds. Every decision compounds without governance.

The Governance Gap

No central registry of agents or their autonomy levels. No reasoning capture for agent decisions. No policy enforcement at the agent layer. When an agent denies a health claim, no one can trace why — not the policyholder, not the insurer, not IRDAI.

THE CORE PROBLEM

This isn't a technology problem. **It's an agent operations problem.** The agents work. The governance doesn't exist.

What Happens When Agents Decide Without Governance

Most insurtech vendors ship a claims model and call it done. Agentic systems — where AI reasons, decides, and acts autonomously across underwriting, claims, and fraud — need a fundamentally different governance architecture. The gap between deployed agents and governed agents is where regulatory and operational risk lives.

Agents Without Operations	Agents with Rotavision
Deploy claims agent, validate accuracy on test set	Every agent registered with autonomy level and risk classification
No reasoning capture — agent denies claim with no trace	Reasoning capture for every claim decision — full audit trail
Manual fraud investigation triggered only after losses mount	Multi-agent fraud detection with cross-validation across signals
Compliance reviewed post-hoc by the audit team	Continuous fairness monitoring for underwriting agents
No central registry of agents or their autonomy levels	Human-in-the-loop controls for high-value and disputed claims
Hope the underwriting agent doesn't discriminate by postcode or gender	IRDAI-ready audit trails from agent registry to decision output

The Indian Bias Taxonomy for Insurance Agent Fairness

When underwriting agents make autonomous decisions, they encode bias at machine speed. Western fairness tools check for race and gender. Indian agents discriminate through proxies that encode centuries of social hierarchy:

BIAS CATEGORY	PROXY VARIABLES AGENTS USE	IMPACT ON INSURANCE DECISIONS
Caste	Surname, pincode, occupation, education tier	Premium pricing and coverage limits
Religion	Name patterns, locality, festival claims patterns	Underwriting approval rates
Region	State of origin, language, migration status	Risk classification variance by geography
Economic Status	Device type, payment method, hospital choice	Claim scrutiny intensity and denial rates
Gender	Occupation, dependents, maternity assumptions	Health coverage exclusions and higher premiums

THE AGENT FAIRNESS PROBLEM

When a human underwriter discriminates, it affects one application. When an agent encodes bias, it affects every decision at production scale. IRDAI expects insurers to demonstrate that AI-driven underwriting is non-discriminatory and explainable.

Agentic AI for India's Insurance Reality

Not generic automation — autonomous agents solving the specific problems Indian insurers and policyholders face every day. Each use case demands agent governance built for the Indian insurance context.

1. Underwriting Agents

Over 650 million Indians lack any form of life or health insurance. For this population, there is no claims history, no policy track record, and no actuarial baseline. Traditional underwriting relies on decades of data that simply doesn't exist. Insurers are deploying autonomous underwriting agents that ingest alternative signals — hospital records, employment data, geographic risk — and price policies without human review.

But when an underwriting agent uses pin code as a proxy for health risk, or prices women's policies higher based on maternity assumptions, the system must catch it. **Vishwas** monitors every underwriting agent decision for fairness across Indian demographic categories, with explainability in 22 languages. **Orchestrate** manages bounded autonomy — the agent prices within guardrails but escalates edge cases to human underwriters.

2. Claims Intelligence Agents

Indian insurers settled over Rs 59,000 crore in claims in FY24. The volume demands automation, but claims are not a single-agent problem. A motor claims agent assesses vehicle damage from photographs — estimating repair costs and determining liability. A health claims agent validates procedures against policy terms, checking for exclusions and waiting periods. A coordinator agent routes decisions, escalates disputes, and triggers payouts.

Each agent reasons independently, but their decisions compound. If the motor damage agent underestimates repair costs, the settlement agent underpays. Every agent's reasoning chain must be captured, every escalation logged, and the entire multi-agent interaction auditable. **Guardian** monitors each claims agent for drift and reliability. **Orchestrate** manages agent composition, policy enforcement, and human-in-the-loop approvals for high-value claims.

3. Fraud Detection Agents

Insurance fraud costs the Indian industry an estimated Rs 30,000–45,000 crore annually. Hospital-insurer collusion rings bill for procedures never performed. Staged accident networks file identical motor claims across insurers. Identity fraud where the same individual holds policies under different names across companies. Multi-agent fraud detection deploys specialised agents for each signal — a claims pattern agent, a provider network agent, a document forensics agent — with a coordinator synthesising signals and assigning fraud risk scores.

Orchestrate manages multi-agent fraud workflows with cross-validation and escalation policies. **Dastavez** provides document AI agents that detect tampering in medical records, repair estimates, and identity documents — with every decision traceable for regulatory and legal proceedings.

THE COMMON THREAD

Every use case requires the same thing: agents that can be **registered, monitored, explained, and bounded**. The use case is specific. The governance architecture is universal.

The Cost Architecture for Claims Agent Operations

Everyone focuses on cost-per-token. The right metric for claims agents is cost-per-decision. And the 10x cost differences come from agent routing architecture, not provider negotiations. The Trust Cascade routes each claims agent decision to the cheapest sufficient intelligence layer.

"~70% of claims decisions can be handled by rules. ~20% by traditional ML. Only ~10% genuinely benefit from agent reasoning. But most deployments route 100% through agents. That's not strategy — that's waste."

Claims Agent Decision Routing: The Trust Cascade

LAYER	VOLUME (10L CLAIMS)	COST/DECISION	MONTHLY COST
L1: Rules Engine (~70%)	7,00,000	Rs 0.001	Rs 700
L2: Statistical ML (~20%)	2,00,000	Rs 0.01	Rs 2,000
L3: Single Agent (~7%)	70,000	Rs 1	Rs 70,000
L4: Multi-Agent Tribunal (~3%)	30,000	Rs 4	Rs 1,20,000
Cascaded Total	10,00,000	Rs 0.19 avg	Rs 1,92,700
Pure Agentic (all LLM)	10,00,000	Rs 3-15	Rs 30L-1.5Cr

The Six Architectural Sins of Claims Agent Deployment

1. Document Firehose

Sending entire 50-page policy documents to the agent when only 3 clauses are relevant. You're paying for tokens the agent ignores.

2. One-Agent-Fits-All

Frontier models for every claim, including simple cashless approvals that a Rs 0.001 rules engine handles perfectly.

3. Retry Spiral

35% of claims involve OCR retries. That's 35% cost overhead plus latency that kills turnaround times.

4. No Semantic Caching

Same hospital bill pattern processed 100 times = 100 identical inference costs. No shared reasoning across claims.

5. No Fraud Pre-Filter

Running expensive fraud analysis on 100% of claims when rules catch 70% of fraud patterns instantly.

6. Verbose Agent Output

Agent asked for approve/reject, responded with three paragraphs. Output tokens cost 3-4x input tokens.

THE MULTIPLIER EFFECT

These sins multiply: 2x (firehose) x 1.5x (one-agent) x 1.35x (retry) x 1.4x (no cache) x 1.5x (verbose) = **8.5x optimal cost**. Agent operations architecture eliminates this waste.

The Regulatory Framework for Insurance Agents

IRDAI's AI/ML Guidelines require insurers to demonstrate that AI-driven decisions are non-discriminatory, explainable, and auditable. When your agents are autonomous, every guideline becomes an agent governance requirement. Bima Sugam and IGMS add policyholder-facing obligations on top.

IRDAI Requirements Mapped to Agent Governance

REQUIREMENT AREA	IRDAI EXPECTATION	AGENT GOVERNANCE REQUIREMENT	ROTAVISION
Fairness	Non-discriminatory underwriting and pricing	Continuous fairness monitoring across Indian demographic categories	Vishwas
Explainability	Decisions interpretable to policyholders	Reasoning capture and vernacular explainability for every agent decision	Vishwas
Accountability	Clear ownership of AI-driven outcomes	Agent registry with ownership, audit trail from registry to output	AgentOps
Governance	Board-level oversight of AI risk	Agent risk classification, autonomy boundaries, periodic reviews	Orchestrate
Fraud Detection	Zero-tolerance fraud monitoring	Multi-agent fraud workflows with cross-validation and escalation	Orchestrate
Data Protection	DPDP Act compliance for policyholder data	Consent capture at agent level, data localization, breach protocols	Dastavez
Grievance Redressal	Bima Sugam / IGMS integration	Agent decision explanations in IGMS-compatible formats	Guardian

Insurance Agent Oversight Structure

IRDAI Agent Governance Structure			
TIER	ROLE	AGENT REQUIREMENTS	CURRENT REALITY
Board	Strategic oversight of agent risk appetite and autonomy boundaries	AI policy approval, agent risk classification, quarterly reviews	Most boards lack agent literacy
Compliance	Independent validation, agent bias audits, IRDAI reporting	Agent challenge sessions, continuous monitoring of agent decisions	Few have dedicated agent oversight
Operations	Build, register, test agents in production	Agent registration, reasoning capture, bounded autonomy config	Basic documentation only

THE COMPLIANCE REALITY

If your agents can't explain their decisions to policyholders, **IRDAI compliance is impossible**. When a claim is denied, the policyholder filing a grievance through Bima Sugam or IGMS has the right to know which agent decided what, and why.

The Agent Operations Stack for Insurance

Deploying a claims agent is not the same as operating one. The Agent Operations Stack is the infrastructure layer between your agents and production — ensuring every agent is registered, governed, monitored, and bounded before it makes a single decision affecting a policyholder.

"The industry doesn't have an agent deployment problem. It has an [agent operations problem](#). The agents work. The infrastructure to govern them doesn't exist."

Five Layers of Agent Operations

1 Agent Registry

Every agent registered with a unique identity, version, owner, risk classification, and autonomy level. No agent operates in production without registration. The single source of truth for what underwriting, claims, and fraud agents exist in your enterprise, what they're authorised to do, and who owns them.

2 Policy Engine

Declarative policies enforced at gateway, sidecar, and inline layers. Policies define what agents can access, what claims they can settle, cost thresholds, and escalation triggers. Policy as code — version-controlled, auditable, and enforceable in real time across underwriting, claims, and fraud workflows.

3 Reasoning Capture

The flight recorder for agent decisions. Every reasoning chain, tool call, intermediate step, and final output captured with full provenance. When IRDAI asks why an agent denied a health claim, you have the complete trace — not a log file, but a reconstructable decision path the policyholder can understand.

4 Bounded Autonomy

Agents decide within guardrails. Low-value cashless approvals are fully autonomous. Medium-risk decisions require post-hoc review. High-value claims and disputed settlements trigger synchronous human-in-the-loop approval. The boundaries are configurable per agent, per line of business, per risk tier.

5 Human-in-the-Loop

Not a checkbox — a workflow. When agents escalate, humans receive the full reasoning chain, the agent's confidence assessment, and the specific policy trigger that caused escalation. Decisions are logged back into the agent's learning loop. Policyholder grievances are traceable to specific agent actions.

THE ROTAVISION DIFFERENCE

Operations, not just deployment. Every layer is [built for regulated insurance](#) — where an ungoverned agent isn't just an engineering risk, it's a policyholder rights violation.

Five Gates for Insurance Agent Production

Before any agent launches in Indian insurance, it must clear five gates. These aren't bureaucratic hurdles — they're the foundations of agent operations that will satisfy IRDAI inspection and protect policyholders at scale.

1 Gate 1: Agent Registration

Agent registered in enterprise registry with unique identity, version, owner, and IRDAI risk classification. Autonomy level defined — fully autonomous, supervised, or human-in-the-loop. Permitted actions, data access, and cost boundaries documented. No unregistered agents in production.

2 Gate 2: Reasoning Capture

Flight recorder active for every agent decision. Complete reasoning chain — inputs, intermediate steps, tool calls, outputs — stored with full provenance. Audit trail reconstructable for any historical claim decision. Retention policy aligned to IRDAI record-keeping requirements.

3 Gate 3: Reliability Monitoring

Drift detection enabled for agent behaviour over time. Hallucination detection active — catching confident wrong answers before they reach policyholders. Sandbagging detection for agents that underperform on edge cases. Alerts configured with on-call routing for production incidents.

4 Gate 4: Fairness and Explainability

Indian bias taxonomy monitoring active — caste proxies, religious inference, regional discrimination, gender, economic status. Explainability in 22 languages for policyholder-facing decisions. IRDAI AI/ML guideline alignment documented. Continuous monitoring, not one-time testing.

5 Gate 5: Bounded Autonomy

Policy enforcement configured and tested. Human-in-the-loop workflows active for high-value and disputed claims. Escalation paths defined and tested — when the agent doesn't know, it asks. Cost controls, rate limits, and budget caps operational. Graceful degradation to lower-cost layers defined.

"An agent should not launch until all five gates are cleared. In Indian insurance, this isn't optional — it's the [minimum bar for agent operations](#) and IRDAI compliance."

Agent Governance Infrastructure for Indian Insurance

Rotavision provides the complete agent governance infrastructure for Indian insurance. Five products built from first principles for agent operations, Indian bias, Indian languages, and IRDAI regulations.

Orchestrate

Multi-Agent Orchestration and Governance

Enterprise-grade agent orchestration with Trust Cascade routing, policy enforcement, and bounded autonomy. Agent registry, reasoning capture, and human-in-the-loop workflows. The operational backbone for governed agent deployment across claims, underwriting, and fraud.

Vishwas

Fairness and Explainability for Agent Decisions

The only fairness system built on the Indian Bias Taxonomy — detecting caste proxies, religious inference, and regional discrimination in underwriting agent decisions. Policyholder-facing explainability in 22 languages. IRDAI AI/ML guideline alignment.

Guardian

Agent Reliability and Drift Monitoring

Continuous production monitoring for claims and underwriting agent behaviour. Catches drift, hallucination, and sandbagging before they impact policyholder decisions. 96% detection accuracy at less than 50ms overhead. IRDAI-compliant documentation generated automatically.

Dastavez

Document AI Agents for Policy and Claims

Multi-script OCR agents for Indian insurance documents — hospital bills in regional formats, handwritten prescriptions, policy documents mixing English with regional scripts. Document tampering detection for fraud investigations. 22-language support.

AgentOps

Enterprise Agent Registry, Policy and Reasoning Capture

From RotaScale. Centralised agent registry with identity, autonomy levels, and risk classification. Declarative policy engine enforced at runtime. Flight recorder for every agent decision. The control plane for enterprise agent operations across the entire insurance value chain — from underwriting through claims to fraud.

BUILT FOR INDIAN INSURANCE. AGENT-FIRST.

Your infrastructure. On-premise, private cloud, or hybrid. **No data leaves India.** Every product built for agent governance in regulated insurance. IRDAI compliant from day one.

What Insurers Build With Agent Governance

Production agent systems processing decisions across underwriting, claims, fraud, and policyholder service. Each implementation demonstrates what becomes possible when agents have proper operations infrastructure.

Fair Underwriting Agent

Autonomous underwriting agent serving 650M+ uninsured populations. Vishwas monitors every decision for caste proxy, religious inference, and regional bias. Full reasoning capture for adverse action explanations. Bounded autonomy with human escalation for edge cases.

Result: 40% faster decisions with zero fair underwriting violations detected

Multi-Agent Claims Intelligence

Motor damage agent, health claims agent, and coordinator agent operating as a governed multi-agent system. Trust Cascade routes 70% through rules, 20% through ML, 10% through agents. Full audit trail for every claim settlement.

Result: 80% straight-through processing, 3-minute turnaround

Multi-Agent Fraud Detection

Claims pattern agent, provider network agent, document forensics agent, and coordinator agent synthesising signals. Hospital-insurer collusion detection. Staged accident ring identification. Evidence chains for legal proceedings.

Result: 94% fraud detection at 86% lower cost than pure agentic approach

Motor Claims Assessment Agent

Image analysis agents for vehicle damage assessment. Integration with garage networks. Fraud detection on repair estimates. Automated settlement for straightforward claims. Human escalation for complex disputes and total loss determinations.

Result: 65% faster settlement, 23% cost reduction

Policyholder Grievance Explainability

Agent decision explanations surfaced through Bima Sugam and IGMS-compatible formats. When a claim is denied, the policyholder receives a vernacular explanation of which agent decided what and why. Full reasoning trace for dispute resolution.

Result: Grievance resolution time reduced by 60%

Enterprise Agent Registry

AgentOps deployed as the central control plane. Every underwriting, claims, and fraud agent registered with identity, autonomy level, and IRDAI risk classification. Policy engine enforces boundaries in real time. Flight recorder captures every decision.

Result: Complete agent inventory with full governance traceability

"The platform doesn't replace your AI strategy — it makes your agents **production-ready for IRDAI regulations**. Same capabilities, but with the governance infrastructure regulators expect."

IRDAI Agent Governance Accelerator

A combined assessment, platform, and integration package for insurers deploying AI agents across underwriting, claims, and fraud — with IRDAI inspection readiness built in.

What's Included

1 Insurance Agent Maturity Assessment

Audit across underwriting, claims, and fraud agent governance against IRDAI AI/ML guidelines. Gap analysis with roadmap for inspection readiness. Board-ready report with prioritised remediation plan.

2 Agent Registry and Policy Engine

Orchestrate + AgentOps configured for insurance — claims agents, underwriting agents, and fraud agents registered with IRDAI risk tiers and autonomy boundaries. Policy enforcement at runtime across all agent workflows.

3 Claims Platform Integration

Pre-built connectors for claims management systems and TPA platforms where claims agents actually operate. Governance layer alongside existing claims infrastructure. Hospital network and IIB integrations supported.

4 Policyholder Grievance Explainability

Agent decision explanations surfaced through Bima Sugam and IGMS-compatible formats. When a claim is denied, the policyholder gets a vernacular explanation of which agent decided what and why. 22 scheduled languages supported.

5 Actuarial Validation Bridge

Connect underwriting agent decisions back to actuarial models. Dashboard showing agent pricing behaviour vs actuarial assumptions, drift from approved pricing bands. Continuous validation for the appointed actuary.

Platform Stack

Agent orchestration Orchestrate	Fairness and explainability Vishwas
Reliability monitoring Guardian	Document AI Dastavez
Agent registry and policy AgentOps (RotaScale)	

Engagement Options

<p>ASSESSMENT</p> <p>Rs 12L</p> <p>2 weeks. Insurance agent maturity audit. IRDAI gap analysis. Indian bias assessment. Board-ready roadmap.</p>	<p>ACCELERATOR</p> <p>Rs 25L</p> <p>4 weeks. Full IRDAI alignment. Agent registry setup. Compliance dashboard. Executive presentation.</p>	<p>PRODUCTION</p> <p>Rs 40L+</p> <p>8-16 weeks. Full platform deployment. Claims platform integration. Team training. Go-live support.</p>
---	---	---

Agents are settling claims for your policyholders. The question is whether anyone is governing them.

650 million Indians have no insurance. Agents will serve them. Rs 59,000 crore in claims need agents that can explain their reasoning. 30% of health claims involve fraud that demands multi-agent intelligence. IRDAI's AI/ML guidelines demand governance that most insurers haven't built. The agents are already deployed. The operations layer is what's missing.

We'd like to show you where you stand. A 30-minute assessment — not a sales pitch — to benchmark your agent governance against IRDAI requirements and identify your highest-value opportunities.

[Request Assessment](#)

Rotavision Consulting Private Limited

HD37, Block D3, Manyata Tech Park,
Outer Ring Road, Venkateshapura,
Bangalore North, Bangalore - 560045, Karnataka, India

CIN: U70200KA2025PTC196547

rotavision.com
hello@rotavision.com
Mumbai | Bengaluru