HEALTHCARE

# Agent Governance for Indian Healthcare

A strategic guide to governing autonomous AI agents in hospitals, health systems, and healthtech. From CDSCO SaMD compliance to clinical agent operations at scale.

EXECUTIVE SUMMARY

Diagnostic agents are reading scans. Triage agents are routing patients. Treatment protocol agents are recommending drugs. India's 1:1,500 doctor-patient ratio is forcing clinical AI adoption faster than governance can keep up. With 80% of specialists concentrated in urban India and 67% of healthcare spending coming out of patients' pockets, an unexplained agent decision isn't a compliance risk — it's a patient safety crisis. CDSCO's SaMD classification framework and the ABDM ecosystem demand agent governance that most health systems haven't built. This guide provides the roadmap.

# Agents Are Making Clinical Decisions. Patient Lives Are at Stake.

Indian healthcare has moved beyond static diagnostic models. Autonomous agents now triage patients, read scans, recommend treatments, and check drug interactions. The shift from model inference to agent autonomy changes the patient safety problem entirely.

**1:1,500**
DOCTOR–PATIENT RATIO — WHO STANDARD IS 1:1,000

**80%**
OF SPECIALISTS CONCENTRATED IN URBAN INDIA

**67%**
OF HEALTHCARE SPENDING IS OUT-OF-POCKET

## Four Gaps That Define the Crisis

### The Access Gap

India has roughly one allopathic doctor for every 1,500 citizens. In rural districts, the ratio can exceed 1:10,000. AI agents are filling the gap by default — triaging patients, reading scans, and recommending treatments with near-zero human review. The scale of the shortage is forcing adoption faster than governance can follow.

### The Specialist Gap

80% of India's specialists are concentrated in urban centres. A patient in a rural PHC in Chhattisgarh with a suspicious chest X-ray has no radiologist within 100 kilometres. Diagnostic imaging agents are stepping in as the first — and often only — reader of clinical scans across thousands of facilities.

### The Governance Gap

No central registry of clinical agents or their risk levels. No reasoning trace for diagnostic decisions — black box diagnosis at scale. When a triage agent routes a chest pain patient to a district hospital 80 kilometres away, no one can reconstruct the reasoning chain. These aren't chatbots — they're clinical decision-makers.

### The Cost Gap

67% of India's healthcare expenditure comes directly from patients' pockets. When an AI agent misdiagnoses or over-prescribes, the financial burden falls on people who can least afford it. Agent errors at scale don't just harm health outcomes — they push families into medical bankruptcy.

THE CORE PROBLEM

This isn't a technology problem. It's a clinical agent operations problem. The agents work. The governance to keep patients safe doesn't exist.

# What Happens When Clinical Agents Operate Without Oversight

Most health-AI vendors ship an agent validated on a curated dataset and call it done. Agentic clinical systems — where AI reasons across symptoms, lab results, and imaging to recommend treatment — need a fundamentally different safety architecture. The gap between deployed agents and governed agents is where patient harm lives.

| Clinical Agents Without Operations | Clinical Agents with Rotavision |
|---|---|
| Deploy diagnostic agent, validate on curated test dataset | Every agent registered with clinical risk classification and autonomy boundaries |
| No central registry of clinical agents or their risk levels | Reasoning capture for every diagnostic recommendation — full audit trail |
| No reasoning trace for clinical decisions — black box diagnosis | Human-in-the-loop mandatory for treatment decisions — bounded autonomy by design |
| Compliance reviewed only after a patient safety incident | Continuous monitoring for demographic and regional bias in clinical outcomes |
| No monitoring for demographic or regional bias in outcomes | Hallucination detection for drug interactions, contraindications, and dosage errors |
| Hope the agent doesn't hallucinate drug interactions or contraindications | ABDM-compliant health data exchange with consent management and complete audit trails |

## The Indian Healthcare Bias Taxonomy

When agents make autonomous clinical decisions, they encode bias at machine speed. Western fairness tools check for race and gender. Indian clinical agents discriminate through proxies that global tooling doesn't detect:

| BIAS CATEGORY | PROXY VARIABLES AGENTS USE | IMPACT ON CLINICAL DECISIONS |
|---|---|---|
| **Caste** | Surname, locality, hospital tier in training data | Different treatment recommendations for same condition |
| **Economic Status** | Insurance type, hospital choice, device ownership | Expensive treatments suggested only for private patients |
| **Rural/Urban** | Location, facility type, imaging equipment | Lower diagnostic confidence for rural imaging data |
| **Skin Tone** | Fitzpatrick IV-VI underrepresentation in training data | Dermatology agents miss conditions on darker skin tones |
| **Gender** | Male-dominated clinical trial data | Drug dosing and side effects miscalculated for women |

THE PATIENT SAFETY PROBLEM

When a human clinician makes a biased decision, it affects one patient. When an agent encodes bias, it affects every patient at production scale. Clinical agent fairness monitoring isn't optional — it's a patient safety obligation.

# Agentic AI for India's Healthcare Reality

Not generic clinical models adapted for India — autonomous agents solving the specific problems India's overstretched healthcare system faces every day. Each use case demands agent governance built for the Indian clinical context.

## 1. Diagnostic Agents

India runs some of the world's largest screening programmes — for tuberculosis, diabetic retinopathy, cervical cancer, and more. The volume of imaging far exceeds available radiologist capacity. Diagnostic agents read chest X-rays in TB screening camps, analyse retinal images for diabetic retinopathy in district hospitals, and flag suspicious findings in CT scans at tertiary centres. These agents process thousands of images daily, often as the first — and sometimes only — reader.

Scale without governance is reckless. A diagnostic agent must flag its confidence level on every read. When confidence falls below clinical thresholds, the case must be escalated to a human radiologist automatically. The agent must be monitored for drift as imaging equipment, patient populations, and disease prevalence change over time. **Guardian** monitors each diagnostic agent for accuracy drift and confidence degradation. **Orchestrate** manages escalation pathways and enforces mandatory human review for uncertain cases.

## 2. Treatment Protocol Agents

India's rural primary health centres serve populations of 20,000–30,000 people, often staffed by a single medical officer with no specialist backup. Treatment protocol agents reason across symptom histories, lab results, and local disease prevalence to recommend treatment pathways, check drug interactions, and suggest dosage adjustments. These agents don't just answer questions — they produce clinical recommendations that doctors act on.

Clinical AI must suggest, not decide. When an agent recommends a treatment protocol, the evidence citations must be verifiable, the confidence level transparent, and the escalation path clear. **Vishwas** ensures every recommendation is explainable — with reasoning traces in 22 Indian languages and evidence citations linked to clinical guidelines. **Orchestrate** enforces bounded autonomy and human-in-the-loop approval for all treatment decisions.

## 3. Triage and Referral Agents

When a triage agent at a PHC in Chhattisgarh routes a patient with chest pain to a district hospital 80 kilometres away, the reasoning must be traceable — in the clinician's language, not in developer logs. Triage agents assess symptom urgency, determine the appropriate level of care, and generate referral recommendations across India's tiered health system. They operate at the point where speed and accuracy directly determine patient outcomes.

Bounded autonomy means the agent operates within clinical guardrails and escalates when it encounters cases outside its training distribution. **Vishwas** monitors triage decisions for demographic and regional bias. **Dastavez** processes referral documents, health records, and consent forms across India's multilingual document ecosystem, linking patient data to ABHA IDs with full consent management.

THE COMMON THREAD

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

# The Cost Architecture for Clinical Agent Operations

Healthcare AI in India must work within the economics of public health — where a chest X-ray costs Rs 50 and a doctor sees 100 patients daily. Everyone focuses on model accuracy. The right metric for clinical agents is cost-per-decision at scale. The Trust Cascade routes each clinical decision to the cheapest sufficient intelligence layer.

"~60% of clinical agent decisions can be handled by rule-based screening. ~25% by specialised ML models. Only ~15% genuinely benefit from agent reasoning. But most deployments route 100% through agents. That's not strategy — that's waste that makes AI inaccessible."

## Clinical Decision Routing: The Trust Cascade

| LAYER | VOLUME (1L CASES) | COST/DECISION | MONTHLY COST |
|---|---|---|---|
| L1: Rule-Based Screening (~60%) | 60,000 | Rs 0.10 | Rs 6,000 |
| L2: Specialised ML Models (~25%) | 25,000 | Rs 1 | Rs 25,000 |
| L3: Single Clinical Agent (~10%) | 10,000 | Rs 10 | Rs 1,00,000 |
| L4: Multi-Agent Clinical Tribunal (~5%) | 5,000 | Rs 50 | Rs 2,50,000 |
| **Cascaded Total** | **1,00,000** | **Rs 3.81 avg** | **Rs 3,81,000** |
| **Pure Agentic (all LLM)** | **1,00,000** | **Rs 25-50** | **Rs 25L-50L** |

## The Six Cost Drivers in Clinical Agent Deployment

### 1. Image Resolution

High-res DICOM files are 50-200 MB each. Transmitting to cloud costs bandwidth. Most screening diagnoses don't need full resolution.

### 2. Model Complexity

Foundation models for every diagnosis is overkill. 60% of chest X-rays are normal — a lightweight screening model handles them at 1/50th the cost.

### 3. Network Dependency

Cloud-first architecture means every scan travels to the cloud and back. In rural India, that's slow, expensive, and often impossible.

### 4. Specialist Escalation

Every AI uncertainty triggers specialist review. Without intelligent confidence thresholds, you're paying specialists to confirm normal findings.

### 5. Redundant Analysis

Same patient, same condition, multiple scans — each analysed independently. Longitudinal context could reduce analysis costs by 30%.

### 6. Verbose Agent Output

Agent asked for a structured finding, responds with three paragraphs. Output tokens cost 3-4x input tokens. Structured extraction cuts costs by 80%.

THE ACCESS MULTIPLIER

Trust Cascade doesn't just cut costs by 10x versus pure foundation models. Edge-first architecture means diagnostics work in rural PHCs without reliable connectivity — reaching the patients who need AI most.

# SaMD Classification for Clinical Agents

CDSCO's Software as Medical Device framework creates a clear classification system for clinical AI agents. When your agents autonomously triage, diagnose, and recommend treatment, every risk class becomes an agent governance requirement. ABDM integration adds a health data governance layer on top.

## CDSCO SaMD Risk Classification for Clinical Agents

**Four-Class Risk Framework for Agent Classification**

| CLASS | RISK LEVEL | CLINICAL AGENT EXAMPLES | GOVERNANCE REQUIREMENT |
|---|---|---|---|
| **Class A** | Low risk | Appointment scheduling agents, general wellness bots | Self-declaration, basic agent registry |
| **Class B** | Low-moderate | Symptom checkers, medication reminder agents | Registration, reasoning capture, basic monitoring |
| **Class C** | Moderate-high | Diagnostic decision support agents, triage agents | Clinical evaluation, bounded autonomy, human-in-the-loop |
| **Class D** | High risk | Cancer detection agents, life-critical diagnostic agents | Clinical trials, full approval, post-market surveillance |

## ABDM Integration Requirements for Clinical Agents

### Health Data Exchange

Any clinical agent accessing patient data must integrate with ABDM health records, use ABHA IDs for identity, and follow consent management protocols. Over 180 million ABHA IDs have been created. Health data is flowing — but who governs the agents consuming it?

### Interoperability Standards

FHIR R4 compliance, SNOMED-CT coding for diagnoses, LOINC for lab results. Clinical agent outputs must be structured for ABDM health information exchange — not just human reading. Consent flows must be auditable at the agent level.

## What's Required Now

### SaMD Classification for Every Clinical Agent

Determine CDSCO risk class (A/B/C/D) for each clinical agent deployed. Diagnostic agents, triage agents, and treatment protocol agents likely fall into Class C or D — requiring clinical evaluation or full approval before production deployment.

### Indian Dataset Validation

CDSCO guidance explicitly requires validation on Indian patient populations. Global approvals alone are insufficient. Clinical agents must demonstrate performance across Indian demographics, skin tones, disease patterns, and imaging equipment variations.

THE COMPLIANCE REALITY

If your clinical agents can't trace their reasoning, CDSCO compliance is impossible. Agent governance isn't a layer you add after deployment. It's the architecture clinical agents must be built on.

# The Safety-First Operations Stack

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

"The industry doesn't have a clinical agent deployment problem. It has a clinical agent operations problem. The agents work. The infrastructure to keep patients safe doesn't exist."

## Five Layers of Clinical Agent Operations

1 **Clinical Agent Registry**

Every agent registered with a unique identity, version, owner, clinical risk classification, and autonomy level. Agents classified by clinical function — triage advisory, diagnostic decision support, treatment recommendation. No agent operates in production without registration and risk-tier assignment.

2 **Clinical Policy Engine**

Declarative policies enforced at gateway, sidecar, and inline layers. Policies define what agents can access, what clinical decisions they can make, escalation triggers, and ABDM data access boundaries. Policy as code — version-controlled, auditable, and enforceable in real time.

3 **Clinical Reasoning Capture**

The flight recorder for clinical agent decisions. Every reasoning chain, evidence citation, intermediate step, and final recommendation captured with full provenance. When CDSCO asks why an agent recommended a treatment, you have the complete trace — not a log file, but a reconstructable clinical decision path.

4 **Bounded Clinical Autonomy**

Agents decide within clinical guardrails. Low-risk screening is fully autonomous. Diagnostic support requires post-hoc clinician review. Treatment recommendations trigger synchronous human-in-the-loop approval. Boundaries are configurable per agent, per use case, per CDSCO risk tier.

5 **Clinical Human-in-the-Loop**

Not a checkbox — a clinical workflow. When agents escalate, clinicians receive the full reasoning chain, the agent's confidence assessment, evidence citations, and the specific policy trigger that caused escalation. Decisions are logged back into the clinical audit trail.

THE ROTAVISION DIFFERENCE

Patient safety, not just deployment. Every layer is built for clinical agent governance — where an ungoverned agent isn't just an engineering risk, it's a patient safety violation.

# Five Gates for Clinical Agent Deployment

Before any clinical agent launches in an Indian hospital or health system, it must clear five gates. In healthcare, these aren't bureaucratic hurdles — they're patient safety obligations. A failed diagnostic in production isn't a bug report. It's a patient harmed.

1 **Gate 1: Agent Registration and Risk Classification**

Agent registered in enterprise registry with unique identity, version, owner, and CDSCO-aligned risk classification. Autonomy level defined — fully autonomous screening, supervised diagnostic support, or human-in-the-loop treatment recommendation. Permitted clinical actions, data access, and ABDM interaction boundaries documented.

2 **Gate 2: Clinical Reasoning Capture**

Flight recorder active for every clinical agent decision. Complete reasoning chain — patient inputs, evidence citations, intermediate steps, confidence scores, final recommendation — stored with full provenance. Audit trail reconstructable for any historical clinical decision. Retention aligned to CDSCO and medico-legal requirements.

3 **Gate 3: Clinical Reliability Monitoring**

Drift detection enabled for diagnostic accuracy over time. Hallucination detection active — catching confident wrong drug interactions and contraindications before they reach clinicians. Performance monitoring across imaging equipment types, patient demographics, and facility tiers. Alerts configured with clinical escalation routing.

4 **Gate 4: Fairness and Clinical Explainability**

Indian healthcare bias taxonomy monitoring active — caste proxies, economic status, rural/urban disparity, skin tone, gender. Clinical explainability in 22 languages for patient-facing and clinician-facing outputs. Evidence grounding against Indian Standard Treatment Guidelines and WHO protocols. Continuous monitoring, not one-time testing.

5 **Gate 5: Bounded Autonomy and Escalation**

Clinical policy enforcement configured and tested. Human-in-the-loop workflows active for treatment decisions and high-risk diagnoses. Escalation paths defined — when the agent encounters cases outside its training distribution, it escalates. Cost controls, rate limits, and graceful degradation to lower-cost screening layers operational.

"A clinical agent should not launch until all five gates are cleared. In Indian healthcare, this isn't optional — it's the minimum bar for patient safety and regulatory compliance."

# Agent Governance Infrastructure for Indian Healthcare

Rotavision provides the complete agent governance infrastructure for Indian healthcare. Five products built from first principles for clinical agent operations, patient safety, Indian regulations, and the reality of healthcare delivery from metro hospitals to district PHCs.

## Orchestrate

**Multi-Agent Orchestration with Human-in-the-Loop Controls**

Enterprise-grade clinical agent orchestration with Trust Cascade routing, policy enforcement, and bounded autonomy. Agent registry with CDSCO-aligned risk classification, clinical reasoning capture, and mandatory human-in-the-loop workflows for treatment decisions. The operational backbone for governed clinical agent deployment.

## Vishwas

**Explainability and Fairness for Clinical Agent Decisions**

The only fairness system built on the Indian Healthcare Bias Taxonomy — detecting caste proxies, economic bias, rural/urban disparity, and skin tone gaps in clinical agent decisions. Patient-facing and clinician-facing explainability in 22 languages. Evidence grounding against Indian Standard Treatment Guidelines.

## Guardian

**Agent Reliability and Drift Monitoring for Diagnostics**

Continuous production monitoring for diagnostic agent accuracy, drift detection across imaging equipment and patient demographics, and adverse event tracking. Catches hallucinated drug interactions and confidence degradation before they impact patient care. CDSCO post-market surveillance ready.

## Dastavez

**Document AI Agents for Health Records and Consent**

Process prescriptions, lab reports, discharge summaries, and consent forms across India's multilingual document ecosystem. Multi-script extraction for ABDM interoperability. FHIR R4 output, ABHA ID linking, and consent management. Works with handwritten documents and low-quality scans common in public health settings.

## AgentOps

**Enterprise Agent Registry, Clinical Risk Classification and Audit**

From RotaScale. Centralised clinical agent registry with identity, autonomy levels, and CDSCO risk-tier classification. Declarative policy engine enforced at runtime. Flight recorder for every clinical agent decision. The control plane for enterprise clinical agent operations.

BUILT FOR INDIAN HEALTHCARE. PATIENT-SAFETY-FIRST.

Your infrastructure. On-premise, private cloud, or hospital data centre. Patient data stays in India. Every product built for clinical agent governance. CDSCO and ABDM compliant from day one.

# What Hospitals and Healthtech Build With Agent Governance

Production clinical agent systems processing decisions across diagnostics, triage, treatment protocols, and patient engagement. Each implementation demonstrates what becomes possible when clinical agents have proper safety operations infrastructure.

### Diagnostic Imaging Agent with Drift Monitoring

Diagnostic agents reading chest X-rays for TB screening across district hospitals. Guardian monitors accuracy drift across equipment types and patient demographics. Trust Cascade routes 60% as normal findings locally, escalates suspicious cases. Validated on Indian patient populations.

`Result: 95% sensitivity, 3-minute turnaround, Rs 3/scan average cost`

### Treatment Protocol Agent with Bounded Autonomy

Autonomous treatment recommendation agent for primary health centres. Vishwas ensures every recommendation is explainable with evidence citations linked to Indian Standard Treatment Guidelines. Orchestrate enforces human-in-the-loop for all treatment decisions. Full reasoning capture for medico-legal audit trail.

`Result: 40% faster treatment decisions with zero guideline deviations`

### Triage Agent with Fairness Monitoring

Triage agents routing patients across the health system based on symptom urgency and facility capability. Vishwas monitors every triage decision for demographic and regional bias. Orchestrate manages escalation pathways with complete audit trails from initial assessment to referral.

`Result: 45% reduction in unnecessary referrals, zero demographic bias detected`

### Vernacular Patient Engagement Agents

Patient-facing agents for medication adherence, symptom monitoring, and follow-up in Hindi, Tamil, Bengali, and 19 other languages. Vishwas monitors every patient interaction for accuracy, safety, and linguistic fidelity. ABDM-compliant with ABHA ID consent management.

`Result: 4.2/5 patient satisfaction, 30% improvement in medication adherence`

### Clinical Documentation Agents

Dastavez agents extracting structured data from handwritten prescriptions and regional-language discharge summaries. FHIR R4 output feeds ABDM health records. Multi-script OCR handles Devanagari, Tamil, Bengali, and Kannada documents from public health facilities.

`Result: 70% reduction in documentation time, 95% extraction accuracy`

### Clinical Agent Registry for Health Systems

AgentOps deployed as the central control plane for a multi-hospital chain. Every clinical agent registered with identity, CDSCO risk classification, and autonomy level. Policy engine enforces clinical boundaries in real time. Flight recorder for every decision across all facilities.

`Result: Complete clinical agent inventory with full governance traceability`

"The platform doesn't replace clinical judgment — it extends expert capability to 10,000 locations where specialists can't be present. Same diagnostic quality, governed by the same safety standards."

# Clinical Agent Safety Accelerator

A combined assessment, platform, and integration package for hospitals and health systems deploying AI agents in clinical workflows — with CDSCO SaMD readiness and ABDM governance built in.

## What's Included

1 **Clinical Agent Risk Classification**

Audit all clinical agents against CDSCO SaMD risk tiers and ABDM data access levels. Map each agent's clinical autonomy boundaries with a readiness roadmap. Identify which agents require Class C evaluation, which require Class D approval, and which can self-declare.

2 **Agent Registry with Clinical Risk Tiers**

Orchestrate + AgentOps configured for healthcare — agents classified by clinical risk (triage advisory vs diagnostic vs treatment), with escalation policies calibrated to risk level. Every agent registered with identity, CDSCO classification, autonomy boundaries, and clinical policy enforcement.

3 **Clinical Evidence Grounding**

Agent outputs grounded against Indian Standard Treatment Guidelines (STGs), National List of Essential Medicines (NLEM), and WHO protocols. Flag when recommendations deviate from approved guidelines. Evidence citation verification for every clinical agent recommendation.

4 **ABDM and HIS Integration**

Pre-built connectors for ABDM Health Data Exchange, ABHA ID consent flows, and major HIS/EMR systems. Agent governance layer deployed alongside clinical workflows. FHIR R4 interoperability for structured clinical agent outputs.

5 **Post-Market Surveillance**

Continuous monitoring of clinical agent performance — outcome tracking, adverse event detection, and demographic disparity analysis. CDSCO-ready vigilance reporting. Guardian deployed for real-time drift detection across all production clinical agents.

## Platform Stack

**Agent orchestration**

Orchestrate

**Clinical explainability**

Vishwas

**Drift monitoring**

Guardian

**Health records AI**

Dastavez

**Agent registry and policy**

AgentOps (RotaScale)

## Engagement Options

ASSESSMENT

### Rs 15L

2 weeks. Clinical agent risk classification. CDSCO SaMD gap

ACCELERATOR

### Rs 30L

4 weeks. Full SaMD alignment. Agent registry setup. Clinical evidence

PRODUCTION

### Rs 50L+

12-20 weeks. Full platform deployment. HIS/EMR integration.

ROTAVISION INDIA

# Clinical agents are making decisions in your hospital. The question is whether anyone is governing them.

A billion Indians need healthcare they can access. With a 1:1,500 doctor-patient ratio and 80% of specialists concentrated in cities, AI agents are stepping in by default. CDSCO's SaMD framework demands governance that most health systems haven't built. The agents are already deployed. The clinical safety 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 clinical agent governance against CDSCO requirements and identify your highest-value patient safety 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
Bengaluru