ENTERPRISE

Enterprise Agent Governance for India

Governing autonomous AI agents across India's largest enterprises — from shadow agent discovery to centralised policy enforcement across every department.

EXECUTIVE SUMMARY

Agent sprawl is the new shadow IT. India's largest enterprises are deploying hundreds of AI agents across departments — credit agents in lending, fraud agents in payments, customer service agents in support, compliance agents in legal. Nobody has a registry. Nobody has a policy engine. Nobody can trace a single decision. The DPDP Act makes the deploying entity accountable, not the vendor. This guide maps the enterprise agent governance journey from L1 (Ad Hoc) to L5 (Autonomous) — centralising discovery, enforcing policy, and building the audit trail before the regulator asks.

Agent Sprawl Is the New Shadow IT.

Every enterprise department is deploying its own agents — credit agents in lending, fraud agents in payments, customer service agents in support, compliance agents in legal. Without a central registry, there's no visibility into what agents exist, what autonomy they have, or what decisions they're making. The governance gap isn't hypothetical — it's already the default state at most large Indian enterprises.

30%+
YOY ENTERPRISE AI SPENDING GROWTH — IDC INDIA, 2025

50+
AI MODELS PER LARGE ENTERPRISE

DPDP
DEPLOYER ACCOUNTABILITY — NOT THE VENDOR

Four Gaps That Define the Problem

Shadow Agent Gap

Departments are spinning up agents independently — procurement agents, HR screening agents, customer service bots, compliance checkers. No central IT team has a complete inventory. When the CISO asks how many agents are active, nobody can answer.

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 rejects a vendor application or flags an employee, no one can trace why.

Accountability Gap

The DPDP Act 2023 makes the deploying entity accountable for AI-driven decisions — not the vendor. If your enterprise's agent denies a service, rejects a claim, or flags a customer, your organisation owns that outcome legally.

Cost Control Gap

Token spend is invisible across departments. No chargeback models. No cost-per-decision tracking. Duplicate agents across teams solving the same problem with different models, different prompts, and no shared infrastructure.

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 Enterprise Agents Go Ungoverned

Departments are buying AI tools. Teams are spinning up agents. But nobody is asking the enterprise-wide question: who registers these agents, who sets their autonomy boundaries, and who captures their reasoning? The gap between deployed agents and governed agents is where operational and regulatory risk lives.

Enterprise Agents Without Operations	Enterprise Agents with Rotavision
Each department deploys its own agents independently	Centralised agent registry with identity, autonomy levels, and risk classification
No central registry or catalogue of active agents	Reasoning capture and flight recorder for every agent
Agent reasoning is a black box across the organisation	Policy enforcement at gateway, sidecar, and inline layers
Compliance team reviews quarterly — at best	Human-in-the-loop controls with bounded autonomy
Shadow agents proliferate across teams	Continuous fairness and reliability monitoring
DPDP accountability unclear when agents make decisions	DPDP and sector-specific regulatory compliance built in

The Agent Maturity Model

Rotavision's Agent Governance Maturity Model maps enterprise readiness from reactive to autonomous. Most Indian enterprises operate at L1 or L2 — agents are deployed but ungoverned.

Enterprise Agent Governance Maturity Model

LEVEL	REGISTRY STATUS	POLICY ENFORCEMENT	REASONING CAPTURE
L1: Ad Hoc	No registry. Agents unknown to central IT.	No policy. Agents operate without boundaries.	No capture. Decisions are untraceable.
L2: Catalogued	Agents listed in spreadsheets. Incomplete.	Guidelines exist. Not enforced at runtime.	Logs exist. Not structured for audit.
L3: Governed	Centralised registry. All agents registered.	Policy enforced at runtime. Per-agent rules.	Flight recorder active. Full audit trail.
L4: Optimised	Registry with drift detection and versioning.	Dynamic policy. Context-aware enforcement.	Continuous monitoring. Anomaly detection.
L5: Autonomous	Self-registering agents with guardrails.	Self-governing within enterprise policy floor.	Real-time reasoning with self-correction.

ENTERPRISE GOVERNANCE AT SCALE

The CIO needs a single pane of glass. Agent governance at enterprise scale means moving from L1 to L3 in months — not years. The maturity model is the roadmap.

Agent Governance Across India's Enterprise Landscape

Every industry deploys agents differently. The governance challenges are sector-specific — shaped by regulators, operational realities, and the stakes of autonomous decisions. The governance architecture is universal.

1. Financial Services Agents

Credit agents are approving loans for populations no model has ever seen. Fraud agents are processing alerts across **21.7 billion monthly UPI transactions**. KYC agents are verifying documents in 12+ Indian scripts. The RBI's FREE-AI framework lays out 7 governance sutras that every regulated entity will need to comply with. Banks and NBFCs deploying agentic AI need agent registries, reasoning capture, and fairness monitoring built for Indian census categories — caste proxies, religious inference, regional discrimination.

Governance requirement: FREE-AI sutra alignment, agent-level audit trails, borrower-facing explainability in 22 languages. **Vishwas** monitors every agent decision for fairness. **Orchestrate** manages bounded autonomy with human-in-the-loop escalation.

2. Telecommunications Agents

Network agents, customer service agents, and fraud detection agents are operating across India's **1.15 billion mobile connections**. Every autonomous decision — from traffic rerouting to complaint resolution in regional languages — needs governance that meets TRAI and DoT requirements. Churn prediction agents are making retention offers. Network optimisation agents are rerouting traffic. Customer service agents are resolving complaints in Hindi, Tamil, Bengali, and 19 other languages.

Governance requirement: TRAI compliance for automated customer interactions, DoT requirements for network-level decisions, DPDP consent management across subscriber base. **Guardian** monitors agent reliability across the network. **Sankalp** ensures data sovereignty.

3. Manufacturing Agents

Quality control agents, predictive maintenance agents, and production optimisation agents are running on **PLI-funded factory floors across 14 sectors**. When an agent rejects a batch or shuts down a production line, the plant manager — and the auditor — needs to know why. Supply chain agents are making procurement decisions. Energy management agents are optimising consumption across facilities.

Governance requirement: Traceability from agent decision to production outcome, safety compliance for autonomous shutdown decisions, PLI audit trails. **AgentOps** provides the central registry. **Guardian** catches drift before it impacts production quality.

THE COMMON THREAD

Every industry deploys agents differently. The governance architecture is universal. Registry, policy, reasoning capture, bounded autonomy, and human-in-the-loop — the same five layers, adapted per sector.

The Cost Architecture for Enterprise Agent Operations

Enterprise AI spending in India is growing over 30% year-on-year. But the cost problem isn't the spend — it's the waste. Most enterprises route 100% of decisions through agents when 70% can be handled by rules. The Trust Cascade routes each decision to the cheapest sufficient intelligence layer.

"~70% of enterprise agent decisions can be handled by rules. ~18% by traditional ML. Only ~12% genuinely benefit from agent reasoning. But most deployments route 100% through agents. That's not strategy — that's waste at enterprise scale."

Enterprise Decision Routing: The Trust Cascade

LAYER	VOLUME (10L)	COST/DECISION	MONTHLY COST
L1: Rules Engine (~70%)	7,00,000	Rs 0.001	Rs 700
L2: Statistical ML (~18%)	1,80,000	Rs 0.01	Rs 1,800
L3: Single Agent (~8%)	80,000	Rs 1	Rs 80,000
L4: Multi-Agent Tribunal (~4%)	40,000	Rs 4	Rs 1,60,000
Cascaded Total	**10,00,000**	**Rs 0.24 avg**	**Rs 2,42,500**
Pure Agentic (all LLM)	**10,00,000**	**Rs 3-15**	**Rs 30L-1.5Cr**

The Six Architectural Sins of Enterprise Agent Deployment

1. Shadow Agents

Departments deploying agents without central visibility. Duplicate agents solving the same problem with different models, prompts, and cost profiles.

2. Duplicate Agents

Three departments running sentiment analysis agents independently. No shared infrastructure. Triple the compute, triple the cost, zero coordination.

3. Ungoverned Spend

No token-level cost tracking. No chargeback models. No department-level budget caps. AI costs buried in cloud infrastructure line items.

4. Frontier for Everything

GPT-5.2 Turbo for classification tasks a rules engine handles perfectly. Enterprise agents defaulting to the most expensive model for every decision.

5. Context Amnesia

No semantic caching across agents. Same query from 50 departments = 50 identical inference costs. No shared reasoning, no shared context.

6. Retry Spiral

35% of enterprise agent requests involve retries. That's 35% cost overhead plus latency that kills internal SLAs and user trust.

THE MULTIPLIER EFFECT

These sins multiply across enterprise scale: 2x (shadow) x 1.5x (duplicate) x 1.35x (retry) x 1.4x (no cache) x 1.5x (frontier) = 8.5x optimal cost. Agent FinOps eliminates this waste.

DPDP and Sector-Specific Agent Governance

The DPDP Act 2023 makes the deploying entity accountable for AI-driven decisions — not the vendor. For enterprises operating across multiple sectors, the regulatory landscape isn't one framework — it's a matrix. Every agent must comply with DPDP at the base layer, plus sector-specific regulations that vary by industry.

DPDP Act 2023: The Enterprise Baseline

Deployer Accountability

If your agent denies a service, rejects a claim, or flags a customer, your enterprise owns that outcome. The vendor that built the model bears no liability. Every agent decision must be traceable to the deploying entity.

Consent and Purpose Limitation

Agents processing personal data require valid consent with clear purpose limitation. When an HR screening agent accesses employee data, the consent framework must cover agentic processing — not just storage.

Sector-Specific Regulation Mapping

REGULATOR	SECTOR	KEY REQUIREMENT	AGENT GOVERNANCE MAPPING
RBI	Banking, NBFCs	FREE-AI framework: 7 Sutras, 26 recommendations	**Vishwas + AgentOps**
IRDAI	Insurance	AI/ML guidelines for underwriting and claims	**Orchestrate + Guardian**
TRAI / DoT	Telecom	Customer protection, network compliance	**Sankalp + AgentOps**
SEBI	Capital Markets	Algorithmic trading and advisory governance	**Guardian + Vishwas**
CDSCO	Healthcare	Medical device and diagnostic AI regulation	**Orchestrate + Vishwas**

Cross-Sector Governance Requirements

Data Localisation

Agent processing and data storage must remain within India. Cross-border data transfers require explicit consent and regulatory approval. Sankalp ensures sovereign routing.

Audit Trail Retention

Agent decision logs must be retained for regulatory periods ranging from 5 to 10 years depending on sector. Flight recorder data must be tamper-proof and reconstructable.

Grievance Redressal

When an agent makes an adverse decision, the affected party must be able to challenge it. The enterprise must produce the agent's reasoning chain in an understandable format.

THE COMPLIANCE REALITY

DPDP makes the deployer accountable, not the vendor. If your agents can't explain their decisions, compliance is impossible. Governance isn't a layer you add later — it's the architecture agents must be built on.

The Enterprise Agent Operations Stack

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

"The CIO needs a single pane of glass for agent governance. Not another dashboard — an operations layer that sits between every agent and every decision."

Five Layers of Enterprise Agent Operations

1 **Agent Registry**

Centralised discovery of all agents across departments with identity and risk classification. Every agent registered with a unique identity, version, owner, department, autonomy level, and data access boundaries. The single source of truth for what agents exist in your enterprise and what they're authorised to do.

2 **Policy Engine**

Enterprise-wide policy enforcement with department-level customisation. Declarative policies enforced at gateway, sidecar, and inline layers. The enterprise sets the policy floor — departments can tighten but never loosen. Policy as code, version-controlled, auditable, and enforceable in real time.

3 **Reasoning Capture**

Flight recorder across every agent in every department. Every reasoning chain, tool call, intermediate step, and final output captured with full provenance. When auditors or regulators ask why any agent made any decision, the complete trace is available — not a log file, but a reconstructable decision path.

4 **Bounded Autonomy**

Configurable per department, per use case, per risk tier. Low-risk departmental decisions are fully autonomous. Medium-risk cross-functional decisions require post-hoc review. High-risk enterprise decisions trigger synchronous human-in-the-loop approval. Boundaries adapt as agents mature.

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

Escalation workflows integrated with existing ITSM — ServiceNow, Jira, PagerDuty. When agents escalate, humans receive the full reasoning chain, confidence assessment, and the specific policy trigger. Decisions are logged back into the agent's learning loop for continuous improvement.

THE ENTERPRISE DIFFERENCE

The CIO needs a single pane of glass for agent governance. Operations, not just deployment. Every layer built for enterprise scale — where ungoverned agents aren't just an engineering risk, they're a regulatory and reputational liability.

Five Gates for Enterprise Agent Production

Before any agent launches across the enterprise, it must clear five gates. These aren't bureaucratic hurdles — they're the foundations of agent operations that satisfy auditors, protect the enterprise, and keep systems reliable at scale across departments.

1 **Gate 1: Agent Registration**

Agent registered in the enterprise registry with unique identity, version, department ownership, and risk classification. Autonomy level defined — fully autonomous, supervised, or human-in-the-loop. Data access boundaries documented. No unregistered agents in production — across any department.

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 from agent decision to regulatory requirement. Retention policy aligned to sector-specific record-keeping mandates.

3 **Gate 3: Reliability Monitoring**

Drift detection enabled for agent behaviour over time across all enterprise agents. Hallucination detection active — catching confident wrong answers before they reach stakeholders. Sandbagging detection for agents that underperform on edge cases. Alerts configured with on-call routing via existing ITSM.

4 **Gate 4: Fairness and Explainability**

Indian bias taxonomy monitoring active for customer-facing agents — caste proxies, religious inference, regional discrimination, gender, economic status. Explainability in the affected party's language, not just the developer's logs. Continuous monitoring, not one-time testing.

5 **Gate 5: Bounded Autonomy**

Department-level guardrails with enterprise-wide policy floor. Human-in-the-loop workflows active for high-stakes decisions. Escalation paths defined and tested. Cost controls, rate limits, and budget caps operational per department. Graceful degradation to lower-cost layers defined.

"An enterprise agent should not launch until all five gates are cleared. In India's regulatory environment, this isn't optional — it's the minimum bar for enterprise agent operations across every sector."

Agent Governance Infrastructure for Indian Enterprises

Rotavision provides the complete agent governance infrastructure for Indian enterprises. Six products built from first principles for agent operations, Indian regulatory compliance, and enterprise-scale deployment across departments.

Sankalp

Sovereign AI Gateway

Route all enterprise agent traffic through India-hosted infrastructure. Data never leaves India. AWS Mumbai, Azure India, private cloud, or air-gapped deployment. DPDP Act compliance by design. The sovereign layer for enterprise agent operations.

Orchestrate

Multi-Agent Orchestration Platform

Enterprise-grade agent orchestration with Trust Cascade routing, cross-department policy enforcement, and bounded autonomy. Manage hundreds of agents across departments from a single control plane. The operational backbone for governed enterprise deployment.

Vishwas

Trust and Fairness Platform

Built on the Indian Bias Taxonomy — detecting caste proxies, religious inference, and regional discrimination in enterprise agent decisions. Explainability in 22 languages for customer-facing agents. DPDP and sector-specific regulatory alignment.

Guardian

AI Reliability Monitoring

Continuous production monitoring for agent behaviour across the enterprise. Catches drift, hallucination, and sandbagging before they impact decisions. 96% detection accuracy at less than 50ms overhead. Integrated with existing observability stacks.

AgentOps

Agent Operations Framework

Centralised agent registry with identity, autonomy levels, and risk classification across the enterprise. Declarative policy engine enforced at runtime. Flight recorder for every agent decision. The control plane for enterprise agent operations at scale.

Dastavez

Document AI Platform

Multi-script OCR agents for Indian documents — Aadhaar, PAN, voter IDs, invoices, contracts in regional formats. Browser agents for web-based document workflows. Intelligent extraction with complete decision audit trails across departments.

BUILT FOR INDIAN ENTERPRISE. AGENT-FIRST.

Your infrastructure. On-premise, private cloud, or hybrid. No data leaves India. Every product built for enterprise agent governance across sectors, departments, and regulatory frameworks.

What Enterprises Build With Agent Governance

Production agent governance systems deployed across India's largest enterprises. Each implementation demonstrates what becomes possible when agents have proper operations infrastructure — from discovery to compliance.

Shadow Agent Discovery and Centralisation

Enterprise-wide audit discovers 200+ ungoverned agents across 12 departments. AgentOps deploys centralised registry with identity, ownership, and risk classification for every agent. Shadow agents either registered and governed or decommissioned. Single source of truth for the CIO.

`Result: Complete agent inventory with 100% registry coverage in 6 weeks`

Enterprise Agent Registry with ITSM Integration

AgentOps integrated with ServiceNow and PagerDuty. Agent incidents routed through existing ITSM workflows. Escalation paths mapped to department ownership. Flight recorder data accessible from ServiceNow tickets for rapid root-cause analysis.

`Result: Agent incident resolution time reduced by 60%`

Cross-Department Agent FinOps

Token consumption tracked per agent, per department, per use case. Trust Cascade routing eliminates frontier model usage for rule-based decisions. Chargeback models allocate AI costs to departments. Budget caps and cost anomaly detection prevent runaway spend.

`Result: 72% reduction in enterprise AI compute costs within 3 months`

DPDP Compliance Automation

Automated consent management at the agent level. Purpose limitation enforcement for every agent processing personal data. Grievance redressal workflows with full reasoning chain transparency. Audit-ready documentation generated on demand for regulatory submissions.

`Result: DPDP audit preparation reduced from weeks to hours`

Multi-Sector Regulatory Dashboard

Real-time compliance monitoring mapped to DPDP, RBI, IRDAI, SEBI, and TRAI requirements. Automated regulatory scoring per agent, per sector. Board-ready governance reports generated on demand. Continuous monitoring dashboards with agent-level drill-down across the enterprise.

`Result: Multi-sector compliance visibility from a single pane of glass`

Agent Maturity Assessment (L1-L5)

Enterprise-wide maturity assessment across Rotavision's Agent Governance Maturity Model. Department-level scoring from L1 (Ad Hoc) to L5 (Autonomous). Prioritised roadmap for moving from current state to target maturity. Quarterly reassessment with continuous improvement tracking.

`Result: Average enterprise maturity improved from L1.2 to L3.1 in 6 months`

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

Enterprise Agent Governance Accelerator

A combined assessment, platform, and integration package for enterprises deploying AI agents across departments — with shadow agent discovery, centralised governance, and DPDP compliance built in.

What's Included

1 **Enterprise Agent Discovery and Audit**

Comprehensive audit of agent deployment across departments — discovery of shadow agents, ungoverned models, and undocumented AI systems. Maturity assessment against Rotavision's Agent Governance Maturity Model (Levels 1–5). Board-ready roadmap.

2 **Centralised Agent Registry and Policy Engine**

Orchestrate + AgentOps deployed enterprise-wide — every agent registered with identity, autonomy level, risk classification, and department ownership. Single source of truth for the CIO. Enterprise policy floor with department-level customisation.

3 **ITSM and Observability Integration**

Pre-built connectors for ServiceNow, Jira, PagerDuty, and existing observability platforms. Agent governance alongside existing IT service management infrastructure. Escalation workflows mapped to department ownership.

4 **Agent FinOps and Cost Governance**

Track token consumption, compute costs, and ROI per agent across the enterprise. Budget allocation, chargeback models, and cost anomaly detection for AI agent operations. Trust Cascade routing for cost optimisation.

5 **DPDP and Sector-Specific Compliance**

Automated compliance mapping across DPDP Act, RBI, IRDAI, SEBI, TRAI, and sector-specific regulations. Audit-ready trails from agent decision to regulatory requirement. Board-ready governance reports on demand.

Platform Stack

Sovereign AI gateway Sankalp	**Agent orchestration** Orchestrate
Fairness and explainability Vishwas	**Reliability monitoring** Guardian
Agent registry and policy AgentOps	**Document AI** Dastavez

Engagement Options

ASSESSMENT

Rs 18L

2 weeks. Shadow agent discovery. Maturity assessment (L1-L5). DPDP gap analysis. Board-ready roadmap.

ACCELERATOR

Rs 35L

6 weeks. Full agent registry. Policy engine deployment. ITSM integration. FinOps setup. Executive presentation.

PRODUCTION

Rs 60L+

12-20 weeks. Full platform deployment. Multi-sector compliance. Team training. Go-live support.

ROTAVISION INDIA

Your enterprise has hundreds of agents. The question is whether anyone knows they exist.

India's largest enterprises are deploying AI agents across every department — without a central registry, without a policy engine, without a single audit trail. The DPDP Act makes you accountable for every decision those agents make. 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 discover your shadow agents, benchmark your governance maturity, and identify your highest-value opportunities.

Request Enterprise 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 | Delhi NCR