MANUFACTURING

Agent Governance for India's Factory Floor

A strategic guide to governing autonomous AI agents in quality control, predictive maintenance, and production optimisation across PLI-funded factories.

EXECUTIVE SUMMARY

Factory agents are making decisions at machine speed — rejecting batches, scheduling maintenance, adjusting production parameters — with cycle times measured in milliseconds. India's manufacturing targets $1 trillion in output with PLI schemes across 14 sectors driving investment. Yet agents controlling quality gates and production lines operate without reasoning capture, safety classification, or ISA-95 alignment. On-premise deployment is non-negotiable; data cannot leave the factory floor. This guide provides the architecture for governed agent operations in Indian manufacturing.

01 The Agent Reality
Factory agents decide. Nobody watches.

02 Agents Without Operations
What ungoverned agents cause

03 Agent Use Cases
Quality, maintenance, production

04 Agent Economics
Cost architecture at scale

05 Regulatory Framework
ISA-95, PLI, BIS compliance

06 Agent Operations Stack
Registry to bounded autonomy

07 Production Readiness
Five gates for factory agents

08 The Platform
Agent governance infrastructure

09 Agent Implementations
What manufacturers build

10 Getting Started
Smart Factory Accelerator

Factory Agents Are Making Decisions. At Machine Speed.

India's manufacturing sector is targeting $1 trillion in output. PLI schemes across 14 sectors are driving massive investment. AI agents on factory floors are now autonomously controlling quality gates, scheduling maintenance, and optimising production parameters — with cycle times measured in milliseconds, not minutes. When an agent rejects a batch or shuts down a line, nobody's capturing why.

$1T+
MANUFACTURING GDP TARGET — MAKE IN INDIA / DPIIT

17%
GDP CONTRIBUTION — GOVERNMENT TARGETS 25% BY 2030

5-10%
PRODUCTION CAPACITY LOST TO UNPLANNED DOWNTIME ANNUALLY

Four Gaps That Define the Problem

The Quality Agent Gap

Vision agents inspect products at line speed — classifying defects, measuring tolerances, rejecting non-conforming parts. But when an agent rejects a batch, the rejection reason is buried in model logits. PLI compliance requires complete quality documentation for every production batch. No reasoning trace means no audit trail.

The Maintenance Agent Gap

Predictive maintenance agents ingest vibration, thermal, and acoustic data from IIoT sensors to predict equipment failures. But when an agent recommends shutting a line, the plant manager gets a probability score — not the sensor evidence, failure mode, or consequences of deferral. With 5-10% capacity lost to downtime, the cost of a wrong call is enormous.

The Governance Gap

No central registry of agents across the plant floor. No reasoning capture for quality rejections or maintenance predictions. No policy enforcement at the agent layer. When a production optimisation agent adjusts parameters autonomously, nobody can reconstruct the decision chain that led to a yield drop.

The Safety Gap

Factory agents interact with physical equipment — adjusting temperatures, speeds, pressures. An agent that controls a furnace setpoint or a robotic arm operates in a safety-critical domain. Without safety integrity classification and bounded autonomy, an ungoverned agent is an operational hazard, not just an engineering risk.

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 Factory Agents Decide Without Governance

Most vendors ship an agent and measure accuracy. Agentic systems on factory floors — where AI reasons, decides, and acts on physical equipment — need a fundamentally different governance architecture. The gap between deployed agents and governed agents is where safety and operational risk lives.

Factory Agents Without Operations	Factory Agents with Rotavision
Deploy quality agent, validate on test images	Every agent registered with autonomy level and blast radius
No reasoning trace for rejection decisions	Reasoning capture for every quality rejection — why was this part rejected?
Maintenance agent runs in a black box	Predictive maintenance agent explains failure predictions with sensor evidence
No central registry of agents across the plant floor	Human-in-the-loop for production line shutdowns
Compliance documentation done quarterly	Continuous drift monitoring for vision agents on changing production batches
Hope the vision model doesn't drift on a new batch	BIS and ISO-compliant audit trails for every agent decision

The Manufacturing Safety Taxonomy

When agents make autonomous decisions on factory floors, safety classification determines governance requirements. Western AI safety frameworks don't account for India's manufacturing conditions — extreme heat, dust, voltage fluctuations, and multi-shift operations:

SAFETY CATEGORY	SAFETY LEVEL	AGENT CONSTRAINTS	GOVERNANCE REQUIREMENT
Equipment Interaction	Critical	No autonomous parameter changes above safety thresholds	Human approval for setpoint changes
Human Safety	Critical	Agent cannot override safety interlocks or lockout/tagout	Full reasoning trace, mandatory escalation
Quality Impact	High	Batch rejection requires documented evidence chain	PLI-compliant audit trail per rejection
Environmental	High	Emissions and discharge within CPCB limits	Continuous monitoring with regulatory alerts
Production Continuity	Medium	Schedule changes within OEE constraints	Post-hoc review, documented rationale

THE SAFETY PROBLEM

When a human operator makes a bad call, it affects one shift. When an agent encodes a bad decision, it affects every production cycle at machine speed. Safety-critical agent governance isn't optional — it's existential.

Agentic AI for India's Factory Floor

Not generic models adapted for India — autonomous agents solving the specific problems Indian manufacturers face every shift. Each use case demands agent governance built for the factory context.

1. Quality Control Agents

Vision agents now inspect products at line speed — classifying defects, measuring tolerances, and rejecting non-conforming parts faster than any human inspector. In PLI-eligible sectors, compliance requires **complete quality documentation for every production batch**. When an agent rejects a batch, the rejection reason must be traceable — not buried in a model's logits.

Agent drift is the hidden risk. A vision model trained on one supplier's raw material degrades silently when the supplier changes, when ambient lighting shifts between seasons, or when a new production batch introduces subtle material variations. Without continuous monitoring, the agent's accuracy erodes — and nobody notices until defective products reach customers or PLI audit documentation falls apart. **Guardian** monitors vision agents for drift in real time. **Vishwas** captures the reasoning behind every rejection decision — traceable, auditable, and ready for BIS and ISO compliance reviews.

2. Predictive Maintenance Agents

India's factories operate in conditions that stress both equipment and models — extreme humidity, pervasive dust, voltage fluctuations, and ambient temperatures that swing 40 degrees between seasons. Predictive maintenance agents ingest vibration data, thermal readings, current draw, and acoustic signatures from **IIoT sensors** to predict failures before they happen.

When an agent recommends shutting a production line, the plant manager needs more than a probability score. They need which sensors triggered the prediction, what failure mode the agent anticipates, and what happens if they defer maintenance by 8 hours. **Guardian** tracks agent prediction accuracy as equipment ages. **Orchestrate** enforces human-in-the-loop policies for line shutdown decisions and captures the full reasoning chain.

3. Production Optimisation Agents

Production optimisation is inherently a multi-agent problem. One agent manages **energy consumption** during peak tariff hours. Another optimises **yield** by adjusting process parameters in real time. A third handles **scheduling** — balancing order priorities, machine availability, and changeover times. These agents must coordinate, not conflict.

For Indian manufacturers, energy costs vary dramatically by state and time of day. PLI compliance demands production documentation demonstrating output targets and quality thresholds are met consistently. **Orchestrate** manages multi-agent coordination across OEE parameters. **Dastavez** generates production documentation automatically from agent decision trails — PLI-ready, BIS-compliant, and audit-proof.

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 Factory Agent Operations

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

"~60% of factory agent decisions can be handled by rules engines. ~25% by traditional ML. Only ~10% genuinely benefit from agent reasoning. But most deployments route 100% through agents. That's not strategy — that's waste."

Agent Decision Routing: The Trust Cascade

LAYER	VOLUME (10L)	COST/DECISION	MONTHLY COST
L1: Rules Engine (~60%)	6,00,000	Rs 0.001	Rs 600
L2: Statistical ML (~25%)	2,50,000	Rs 0.01	Rs 2,500
L3: Single Agent (~10%)	1,00,000	Rs 1	Rs 1,00,000
L4: Multi-Agent Tribunal (~5%)	50,000	Rs 4	Rs 2,00,000
Cascaded Total	**10,00,000**	**Rs 0.30 avg**	**Rs 3,03,100**
Pure Agentic (all LLM)	**10,00,000**	**Rs 3-15**	**Rs 30L-1.5Cr**

The Six Architectural Sins of Factory Agent Deployment

1. Monolithic Prompts

2,000 tokens of production context for every agent call, including equipment specs the agent ignores. You're paying for tokens that add no value.

2. Sensor Firehose

Stuffing all IIoT sensor data into agent context. Your agent ingests 50 sensor streams when the anomaly is in one vibration channel.

3. Retry Spiral

35% of agent requests involve retries. That's 35% cost overhead plus latency that kills real-time production SLAs.

4. Context Amnesia

No semantic caching. Same quality check from 10 identical parts = 10 identical inference costs. No shared reasoning across production runs.

5. One-Agent-Fits-All

Frontier models for simple threshold checks a rules engine handles perfectly. Using GPT-5 to check if a temperature exceeds a setpoint.

6. Verbose Agent Output

Agent asked for pass/fail, responded with three paragraphs of analysis. Output tokens cost 3-4x input tokens.

THE MULTIPLIER EFFECT

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

ISA-95, PLI, and BIS Compliance for Factory Agents

Factory agent governance doesn't operate in a regulatory vacuum. ISA-95 defines the automation hierarchy. PLI schemes demand auditable production documentation across 14 sectors. BIS and ISO standards govern quality systems. When your agents are autonomous, every standard becomes an agent governance requirement.

ISA-95 / Purdue Model: Agent Governance by Zone

ISA-95 Agent Governance Structure

ISA-95 LEVEL	FUNCTION	AGENT GOVERNANCE REQUIREMENT	ROTAVISION
Level 4: ERP	Business planning, supply chain, order management	Agent decisions auditable against PLI targets and business KPIs	Dastavez + AgentOps
Level 3: MES	Production scheduling, quality management, batch tracking	Agent registry with safety integrity levels, reasoning capture for quality decisions	Orchestrate + Vishwas
Level 2: SCADA	Supervisory control, HMI, real-time monitoring	Bounded autonomy for parameter changes, human-in-the-loop for setpoint overrides	Orchestrate + Guardian
Level 1: PLC	Direct equipment control, sensors, actuators	Agent cannot bypass safety interlocks; all control actions logged with full provenance	Guardian + Sankalp

PLI and Standards Compliance Mapping

PLI Compliance Across 14 Sectors

PLI schemes have attracted Rs 1.03 lakh crore in investment. Every PLI-eligible manufacturer must demonstrate output targets and quality thresholds are met consistently. Agent decision trails must map to DPIIT documentation requirements — production batch reports with full reasoning capture for every quality gate.

BIS and ISO Standards for Agents

ISO 9001 quality management, ISO 55001 asset management, and BIS product standards all require documented decision rationale. When agents make quality and maintenance decisions, the audit trail must satisfy external assessors — not just internal dashboards. Agent reasoning must be reconstructable.

THE COMPLIANCE REALITY

Compliance isn't documentation you write after the fact. It's architecture you build into the agent from day one. ISA-95 alignment, PLI readiness, and BIS compliance are governance requirements, not afterthoughts.

The Agent Operations Stack for Manufacturing

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

"The industry doesn't have an agent deployment problem. It has an agent operations problem. The agents work. The infrastructure to govern them on factory floors doesn't exist."

Five Layers of Agent Operations

1 **Agent Registry**

Every factory agent registered with a unique identity, safety integrity level, OT zone boundary, and permitted equipment interactions. No agent operates on the plant floor without registration. The single source of truth for what agents control, what they're authorised to do, and who owns them.

2 **Policy Engine**

ISA-95 zone-aware policy enforcement at gateway, sidecar, and inline layers. Policies define what equipment agents can interact with, what parameters they can modify, safety thresholds, and escalation triggers. Policy as code — version-controlled, auditable, and enforceable in real time across the OT network.

3 **Reasoning Capture**

The flight recorder for factory agent decisions. Every quality rejection, maintenance prediction, and parameter adjustment captured with full sensor evidence and provenance. When a PLI auditor asks why an agent rejected a batch, you have the complete trace — not a log file, but a reconstructable decision path with sensor data.

4 **Bounded Autonomy**

Safety-critical decisions require human-in-the-loop. Advisory agents operate fully autonomously — recommendations only. Supervised agents adjust parameters within defined ranges. Autonomous equipment control requires synchronous human approval above safety thresholds. Boundaries configurable per agent, per equipment, per safety level.

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

When an agent recommends a line shutdown, the plant manager receives the full reasoning chain — which sensors triggered the prediction, what failure mode is anticipated, and consequences of deferral. Decisions are logged back into the agent's operational record with the manager's rationale for acceptance or override.

THE ROTAVISION DIFFERENCE

Operations, not just deployment. Every layer is built for safety-critical manufacturing — where an ungoverned agent isn't just an engineering risk, it's an operational hazard.

Five Gates for Factory Agent Production

Before any agent launches on an Indian factory floor, it must clear five gates. These aren't bureaucratic hurdles — they're the foundations of agent operations that will satisfy PLI auditors, meet BIS requirements, and keep your production lines safe at scale.

1 **Gate 1: Agent Registration**

Agent registered in enterprise registry with unique identity, safety integrity level, OT zone, and permitted equipment interactions. Autonomy level defined — advisory (no physical impact), supervised (bounded parameter changes), or autonomous (direct equipment control). No unregistered agents on the production floor.

2 **Gate 2: Reasoning Capture**

Flight recorder active for every agent decision. Every quality rejection includes the defect classification, confidence score, and image evidence. Every maintenance prediction includes sensor readings, failure mode analysis, and recommended action. Retention policy aligned to BIS and PLI record-keeping requirements.

3 **Gate 3: Reliability Monitoring**

Drift detection enabled for vision agents across production batches — catching accuracy degradation when suppliers change, lighting shifts, or material variations appear. Sensor model calibration monitoring for predictive maintenance agents. Alerts configured with on-call routing for production incidents.

4 **Gate 4: Safety Classification**

Every agent classified by safety impact: advisory (no physical impact, recommendations only), supervised (parameter adjustments within defined ranges with post-hoc review), or autonomous (direct equipment control with mandatory human-in-the-loop above safety thresholds). Classification determines governance requirements.

5 **Gate 5: Bounded Autonomy**

Human-in-the-loop enforcement active for line shutdowns, equipment parameter changes above safety thresholds, and production schedule overrides. Escalation paths defined and tested — when the agent doesn't know, it asks the plant manager. Cost controls, rate limits, and safety caps operational. Graceful degradation to manual control defined.

"An agent should not launch until all five gates are cleared. In Indian manufacturing, this isn't optional — it's the minimum bar for safe agent operations on the factory floor."

Agent Governance Infrastructure for Indian Manufacturing

Rotavision provides the complete agent governance infrastructure for Indian manufacturing. Six products built from first principles for factory agent operations, ISA-95 alignment, PLI compliance, and on-premise deployment.

Orchestrate

Multi-Agent Factory Operations

Enterprise-grade agent orchestration with Trust Cascade routing, ISA-95 zone-aware policy enforcement, and bounded autonomy for safety-critical decisions. Agent registry, reasoning capture, and human-in-the-loop workflows for plant managers. The operational backbone for governed agent deployment on factory floors.

Guardian

Agent Reliability for Production

Continuous production monitoring for factory agent behaviour. Catches vision agent drift across production batches, sensor model degradation as equipment ages, and prediction accuracy erosion. 96% detection accuracy at less than 50ms overhead — fast enough for real-time production monitoring.

Dastavez

Document AI for Quality and Compliance

Automated production documentation from agent decision trails. PLI batch reports, BIS quality records, and ISO audit documentation generated automatically. Quality certificates, inspection reports, and compliance filings traced back to the agent decisions that produced them.

Vishwas

Agent Decision Explainability

Every quality rejection, maintenance prediction, and production adjustment explained with full reasoning chains. Plant managers see sensor evidence, decision logic, and alternatives considered. Auditor-ready explanations that satisfy PLI reviewers and BIS assessors.

AgentOps

Enterprise Agent Registry and Policy Engine

From RotaScale. Centralised agent registry with safety integrity levels, OT zone boundaries, and equipment interaction permissions. Declarative policy engine enforced at runtime across the industrial network. Flight recorder for every factory agent decision.

Sankalp

Sovereign Deployment, On-Premise Support

Deploy the entire agent governance stack on-premise, within the factory network. Air-gapped deployment for classified manufacturing. Data never leaves the factory floor. ISA-95 network segmentation respected. Private cloud, on-premise, or hybrid — your infrastructure, your control.

BUILT FOR INDIAN MANUFACTURING. AGENT-FIRST.

On-premise. Air-gapped. ISA-95 aligned. No data leaves the factory. Every product built for agent governance in safety-critical manufacturing. PLI-compliant from day one.

What Manufacturers Build With Agent Governance

Production agent systems governing decisions across quality, maintenance, production, compliance, and energy. Each implementation demonstrates what becomes possible when factory agents have proper operations infrastructure.

Vision Quality Agent with Drift Monitoring

Autonomous vision agent inspecting products at line speed with Guardian monitoring for drift across production batches. Vishwas captures reasoning for every reject decision with defect classification and image evidence. Continuous accuracy tracking as suppliers and materials change.

`Result: 92% defect detection with zero unexplained rejections`

Predictive Maintenance Agent with Sensor Evidence

IIoT-connected maintenance agent monitoring vibration, thermal, and acoustic signatures. Every prediction includes sensor evidence chain, failure mode analysis, and deferral risk assessment. Human-in-the-loop for shutdown recommendations via Orchestrate.

`Result: 40% reduction in unplanned downtime with full audit trail`

Multi-Agent OEE Optimisation

Coordinated agents for energy, yield, and scheduling operating across OEE parameters. Orchestrate prevents conflicting optimisations — energy agent can't override yield targets, scheduling agent respects changeover constraints. Every parameter adjustment documented.

`Result: 8% OEE improvement with zero agent conflicts`

PLI Compliance Documentation Agent

Dastavez generates PLI-ready production batch reports from agent decision trails. Output targets, quality thresholds, and compliance metrics mapped to DPIIT documentation requirements across 14 sectors. Audit preparation automated from agent reasoning capture.

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

Energy Cost Optimisation Agent

Production scheduling agent that optimises against state-specific industrial tariff structures and time-of-day pricing. Shifts energy-intensive operations to off-peak hours while maintaining output targets. Every scheduling decision documented with cost-benefit reasoning.

`Result: 12% energy cost reduction with maintained production targets`

Agent Registry for the Factory Floor

AgentOps deployed as the central control plane across all factory agents. Every agent registered with safety integrity level, OT zone, and equipment permissions. Policy engine enforces ISA-95 zone boundaries. Flight recorder captures every decision for compliance audit.

`Result: Complete agent inventory with full governance traceability`

"The platform doesn't replace your automation strategy — it makes your factory agents production-ready for Indian compliance requirements. Same capabilities, but with the governance infrastructure auditors expect."

Smart Factory Agent Governance Accelerator

A combined assessment, platform, and integration package for manufacturers deploying AI agents across quality, maintenance, and production — with PLI audit readiness and ISA-95 alignment built in.

What's Included

1 **Smart Factory Agent Maturity Assessment**

Audit agent readiness against ISA-95/Purdue Model levels. Gap analysis across quality, maintenance, and production agents with PLI compliance readiness roadmap and safety integrity classification for every deployed agent.

2 **Agent Registry and OT Security Integration**

Orchestrate + AgentOps configured for manufacturing — quality agents, maintenance agents, and production agents registered with safety integrity levels and OT zone boundaries. ISA-95 network segmentation respected.

3 **SCADA/MES/ERP Integration Layer**

Pre-built connectors for SCADA, MES, and ERP systems (SAP, Oracle) where factory agents operate. Governance layer alongside existing industrial automation infrastructure. No disruption to production systems.

4 **PLI Compliance Automation**

Agent decision trails mapped to PLI documentation requirements across 14 sectors. Automated production batch reporting with full agent reasoning capture for DPIIT audit readiness. Quality gate documentation generated automatically.

5 **Safety Integrity and Bounded Autonomy**

Safety classification for every factory agent — from advisory (no physical impact) to autonomous (direct equipment control). Human-in-the-loop enforcement for safety-critical line shutdown decisions. Escalation workflows tested before go-live.

Platform Stack

Agent orchestration Orchestrate	**Decision explainability** Vishwas
Reliability monitoring Guardian	**Document AI** Dastavez
Agent registry and policy AgentOps (RotaScale)	**Sovereign deployment** Sankalp

Engagement Options

ASSESSMENT

Rs 15L

2 weeks. ISA-95 gap analysis. Agent safety audit. PLI readiness assessment. Board-ready roadmap.

ACCELERATOR

Rs 30L

4 weeks. Full ISA-95 alignment. Agent registry setup. SCADA/MES integration. Safety classification.

PRODUCTION

Rs 50L+

8-16 weeks. Full platform deployment. ERP integration. On-premise setup. Team training. Go-live support.

ROTAVISION INDIA

India is becoming the world's factory. The agents running those factories need governance — not just uptime.

$1 trillion in manufacturing output. 14 PLI sectors driving investment. Agents controlling quality gates, predicting equipment failures, and optimising production lines at machine speed. The agents are already deployed. The operations layer is what's missing. When an agent shuts down a production line, the plant manager deserves to know why.

We'd like to show you where you stand. A 30-minute assessment — not a sales pitch — to benchmark your factory agent governance against ISA-95 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
Pune | Chennai | Bengaluru