

TELECOMMUNICATIONS

Agent Governance for Indian Telecom

A strategic guide to governing autonomous network agents on the road to AN Level 4. From TM Forum compliance to production agent operations across 1.15 billion connections.

EXECUTIVE SUMMARY

India's telcos are racing to Autonomous Network Level 4 — where AI agents make network decisions independent of humans. Jio's Open Telecom AI Platform is built on agentic AI. Airtel has deployed zero-touch maintenance with Nokia. JioBrain runs digital twins for capacity planning. But 1.15 billion mobile connections are governed by network agents that have no registry, no reasoning capture, and no blast radius controls. AN Level 4 without agent governance isn't autonomy — it's liability. This guide provides the roadmap for closing that gap.

- | | |
|--|---|
| 01 The AN Reality
Racing to Level 4 without governance | 02 Autonomy Without Governance
What ungoverned agents cause |
| 03 Agent Use Cases
Network, digital twins, fraud | 04 Agent Economics
Cost architecture at scale |
| 05 TRAI & DoT Compliance
Regulatory framework for AN | 06 Agent Operations Stack
Registry to bounded autonomy |
| 07 Production Readiness
Five gates for network agents | 08 The Platform
Agent governance for telecom |
| 09 Agent Implementations
What telcos build | 10 AN Governance Accelerator
The solution package |

Racing to Level 4. Without a Governance Framework.

The telecom industry has defined six levels of network autonomy (AN Level 0–5), and 70+ operators globally have committed to reaching Level 4 — where agents make network decisions independent of humans. India's operators are leading this race. The shift from network automation to agent autonomy changes the governance problem entirely.

1.15B

MOBILE CONNECTIONS IN INDIA – TRAI, Q2 FY25

850M+

INTERNET SUBSCRIBERS – TRAI, 2025

L4

AN LEVEL TARGET FOR 2025–27 – TM FORUM

Four Gaps That Define the Problem

The Autonomy Gap

AN Level 4 is the inflection point — where agents decide independently across multiple network domains. Jio's Open Telecom AI Platform with AMD, Cisco, and Nokia is built on agentic AI. Airtel has deployed zero-touch maintenance with Nokia. But autonomy without governance means agents reroute traffic, reallocate spectrum, and shut down cell sectors with no reasoning trace.

The Scale Gap

Every agent decision — traffic routing, slice allocation, fault prediction — propagates across 1.15 billion connections instantly. A misconfigured network agent doesn't affect one subscriber. It cascades network-wide. At this scale, a single-agent error can cause outages affecting millions. Blast radius controls don't exist in most NOCs.

The Governance Gap

No central registry of network agents or their autonomy levels. No reasoning capture for agent decisions. Nobody in the NOC knows how many agents are running, what they're authorised to do, or who owns them. When an agent deprioritises enterprise traffic during peak hours, nobody can trace why.

The Digital Twin Gap

JioBrain runs digital twins for capacity planning. Vodafone Idea adopted HCL's ANA platform for self-healing simulations. But most operators govern twin agents differently from production agents — or don't govern them at all. An agent validated in a twin with no policy enforcement behaves differently when policies are applied in production.

THE CORE PROBLEM

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

What Happens When Network Agents Decide Without Governance

AN Level 4 isn't just more automation — it's the point where agents make decisions independent of humans. That transition requires governed autonomy: agent registration, reasoning capture, policy enforcement, and blast radius controls. Without it, you're not autonomous — you're uncontrolled.

Autonomy Without Governance	Governed Autonomy with Rotavision
Deploy network agents, validate on test traffic	Every agent registered with autonomy level, blast radius, and risk classification
No central registry — nobody knows how many agents run in the NOC	Reasoning capture for every network decision — traceable from agent to outcome
Agent reroutes traffic or reallocates spectrum with no reasoning trace	Digital twin agents governed with the same policy engine as production agents
Digital twin runs simulations but agent decisions aren't captured	Blast radius controls that prevent single-agent cascades across a billion connections
No blast radius controls — one agent misconfiguration cascades network-wide	Policy enforcement at gateway and inline layers — bounded autonomy by design
AN Level self-assessment is a PowerPoint exercise	AN Level readiness mapped to agent maturity, not just network capability

The Telecom Agent Risk Taxonomy

When network agents make autonomous decisions, they create risk at machine speed. A fault prediction agent that shuts down a healthy cell sector. A slice management agent that deprioritises emergency services. A capacity agent that underserves rural circles. These risks compound across 1.15 billion connections:

RISK CATEGORY	AGENT BEHAVIOUR	IMPACT AT SCALE
Cascade Failure	Agent reroutes traffic without blast radius limits	Single misconfiguration affects millions of subscribers
Service Equity	Capacity agents optimise for ARPU, not coverage	Rural circles and low-revenue areas systematically underserved
Spectrum Waste	Slice agents over-provision without cost awareness	Spectrum resources consumed with no corresponding revenue
Regulatory Exposure	Agents violate QoS mandates with no audit trail	TRAI penalties and license review for persistent non-compliance
Twin Divergence	Digital twin agents behave differently from production	Validation in twin provides false confidence for production

THE AGENT GOVERNANCE PROBLEM

When a human NOC engineer misconfigures a router, it affects one circuit. When a network agent encodes a bad policy, **it affects every decision across 1.15 billion connections**. Agent governance isn't optional — it's existential.

The Path to Autonomous Networks

Not generic network automation adapted for India — governed agent systems that move Indian telcos from Level 2 to Level 4 and beyond. Each use case demands agent governance built for the scale and complexity of India's network infrastructure.

1. The Road to AN Level 4

AN Level 4 is the inflection point — where agents make network decisions independent of humans across multiple domains. For India's operators, this means agents that predict faults across **1.15 billion connections**, trigger self-healing workflows, manage 5G network slice allocation in real time, and optimise capacity across terrain from Himalayan base stations at 4,500 metres to ultra-dense urban cells in Mumbai. Jio is building towards this with its **Open Telecom AI Platform**. Airtel has deployed **zero-touch maintenance with Nokia**.

But Level 4 autonomy requires Level 4 governance. When a fault prediction agent shuts down a cell sector, the reasoning chain must be traceable. When a slice management agent deprioritises enterprise traffic during peak hours, the decision must be explainable. **Orchestrate** manages the agent fleet across the NOC — registration, autonomy levels, policy enforcement, and closed-loop control. **Guardian** monitors every network agent for drift and reliability.

2. Network Digital Twins

JioBrain operates a digital twin system for capacity planning and early failure warnings. Vodafone Idea adopted HCL's ANA platform for self-healing simulations. The principle is sound: simulate agent behaviour in a twin before production. But here's the gap — most operators govern twin agents differently from production agents, or don't govern them at all. An agent validated in a twin with no policy enforcement behaves differently when policies apply in production.

The twin must mirror production governance exactly — same agent registry, same policy engine, same reasoning capture. **Guardian** applies identical monitoring to agents in both environments. **Orchestrate** ensures the same policy engine governs agents wherever they run.

3. Fraud and Customer Intelligence

Telecom fraud in India operates across multiple vectors simultaneously. SIM swap fraud exploits Aadhaar-linked verification, CLI spoofing impersonates legitimate numbers, bypass fraud routes international calls through local gateways. Airtel's AI-driven spam detection has **reduced subscriber complaints by 49%**. Meanwhile, TRAI mandates quality-of-service standards requiring customer service agents to handle 22 languages with traceable decisions.

Multi-agent systems outperform monolithic models. A SIM swap pattern agent, a CLI verification agent, and a call-flow analysis agent each reason independently, then a coordinator synthesises signals. **Vishwas** monitors customer-facing decisions for fairness. **Orchestrate** manages multi-agent fraud workflows with complete audit trails.

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 Network Agent Operations

Everyone focuses on cost-per-token. The right metric for network 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.

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

Agent Decision Routing: The Trust Cascade

LAYER	VOLUME (10L DECISIONS)	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 1.5	Rs 1,05,000
L4: Multi-Agent Tribunal (~3%)	30,000	Rs 5	Rs 1,50,000
Cascaded Total	10,00,000	Rs 0.26 avg	Rs 2,57,700
Pure Agentic (all LLM)	10,00,000	Rs 4-18	Rs 40L-1.8Cr

The Six Architectural Sins of Network Agent Deployment

1. Every Decision to LLM

80% of network decisions are routine — threshold alerts, capacity checks, standard reroutes. Sending these through agent reasoning is burning money on solved problems.

2. No Domain Routing

Using a single general-purpose agent for radio, transport, and core decisions. Specialised agents per domain are 8x cheaper and more accurate for specific network operations.

3. Full Context Always

Sending complete network topology on every agent call. 95% of decisions need only local context. You're paying for tokens the agent ignores.

4. No Semantic Caching

Same fault pattern from 500 cell sites = 500 identical inference costs. Semantic caching cuts this to near zero for repeated patterns.

5. One-Agent-Fits-All

Frontier models for anomaly detection that statistical methods handle at 1/100th the cost with equal accuracy. Right-size intelligence to the decision.

6. Retry Spirals

Network latency causes agent timeouts. Timeouts trigger retries. Retries multiply costs. 30% overhead from poor timeout handling in NOC environments.

THE MULTIPLIER EFFECT

These sins multiply: 2x (every decision) x 1.5x (no routing) x 1.4x (full context) x 1.5x (no cache) x 1.3x (retry) = **8.2x optimal cost**. Agent operations architecture eliminates this waste.

The Regulatory Framework for Autonomous Networks

Telecom is designated strategic national infrastructure. As operators deploy autonomous network agents, they face a layered regulatory framework: TRAI's quality-of-service mandates, DoT's spam coordination requirements, AI calling regulations with penalties up to Rs 10 lakh, and the emerging requirement to demonstrate agent governance to both regulators and the TM Forum AN assessment framework.

Regulatory Requirements Mapped to Agent Governance

REGULATION	REQUIREMENT	AGENT GOVERNANCE IMPLICATION	ROTAVISION
TRAI QoS	Service quality standards across all circles	Agents must not systematically underserve regions or segments	Vishwas
AI Calling	Mandatory AI disclosure, consent registry	Customer-facing agents must comply with disclosure and consent	Orchestrate
DoT Spam	Spam detection, cross-operator coordination	Fraud agents must integrate with DoT framework, audit trails required	Guardian
DPDP Act	Data protection, consent, purpose limitation	Agent data access governed by purpose, subscriber consent captured	Sankalp
TM Forum AN	AN Level assessment and maturity framework	Agent governance maturity maps directly to AN Level readiness	AgentOps

AN Level Governance Requirements

Agent Governance by AN Level			
AN LEVEL	AGENT AUTONOMY	GOVERNANCE REQUIRED	INDUSTRY STATUS
L0-L1	Manual / Assisted — humans decide, agents suggest	Basic logging and decision support documentation	Legacy — most operators have moved beyond
L2-L3	Partial / Conditional — agents decide within narrow domains	Agent registry, reasoning capture, domain-specific policies	Where most Indian operators are today
L4	Highly Autonomous — agents decide across domains independently	Full governance: registry, policy engine, blast radius, audit	Industry target 2025-27 — governance gap is critical
L5	Full Zero-Touch — no human involvement	Self-governing agents with provable safety guarantees	Future — requires governance breakthroughs

THE COMPLIANCE REALITY

If your network agents can't explain their decisions, **AN Level 4 compliance is impossible**. Governance isn't a layer you add after autonomy. It's the architecture agents must be built on.

The Agent Operations Stack for Telecom

Deploying a network 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 in your network.

"The industry doesn't have a network automation 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 network agent registered with a unique identity, version, owner, risk classification, and autonomy level. No agent operates in the NOC without registration. The single source of truth for what agents exist in your network, what they're authorised to do, and who owns them. Maps directly to TM Forum AN Level assessment requirements.

2 Policy Engine

Declarative policies enforced at gateway, sidecar, and inline layers. Policies define what agents can access, what network decisions they can make, blast radius limits, and escalation triggers. Intent-to-policy translation converts business intent ("prioritise video in Mumbai during IPL") into agent guardrails. Policy as code — version-controlled and enforceable in real time.

3 Reasoning Capture

The flight recorder for network agent decisions. Every reasoning chain, tool call, intermediate step, and final output captured with full provenance. When TRAI asks why an agent deprioritised a cell sector, you have the complete trace — not a log file, but a reconstructable decision path.

4 Blast Radius Controls

Agents decide within guardrails that limit the scope of impact. Low-risk decisions are fully autonomous. High-impact decisions — those affecting capacity across circles or triggering service degradation — require synchronous human-in-the-loop approval. Boundaries are configurable per agent, per domain, per risk tier. A single-agent error can't cascade across a billion connections.

5 Human-in-the-Loop

Not a checkbox — a workflow. When agents escalate, NOC engineers 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. The human doesn't slow the network — they govern the boundary where agent autonomy meets network safety.

THE ROTAVISION DIFFERENCE

Operations, not just deployment. Every layer is [built for telecom at billion-subscriber scale](#) — where an ungoverned agent isn't just an engineering risk, it's a network-wide incident.

Five Gates for Network Agent Production

Before any network agent launches in Indian telecom, it must clear five gates. These aren't bureaucratic hurdles — they're the foundations of agent operations that will satisfy TRAI auditors, pass TM Forum AN assessments, and keep your network reliable at billion-subscriber scale.

1 Gate 1: Agent Registration

Agent registered in enterprise registry with unique identity, version, owner, and risk classification. Autonomy level defined — fully autonomous, supervised, or human-in-the-loop. Blast radius limits documented. Permitted network domains, actions, and escalation triggers specified. No unregistered agents in the NOC.

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 network decision. Retention aligned to TRAI record-keeping and TM Forum assessment requirements.

3 Gate 3: Reliability Monitoring

Drift detection enabled for agent behaviour over time. Hallucination detection active — catching confident wrong decisions before they affect network operations. Twin-production divergence monitoring ensures validated behaviour matches production reality. Alerts configured with NOC routing for production incidents.

4 Gate 4: Service Equity and Compliance

Fairness monitoring active across circles, subscriber segments, and geographies. Agents must not systematically underserve rural areas or low-ARPU subscribers. TRAI QoS compliance validated. AI calling disclosure and consent registry integration tested for customer-facing agents. Regulatory documentation audit-ready.

5 Gate 5: Blast Radius and Cost Controls

Blast radius limits configured and tested. Policy enforcement active for high-impact decisions. Cost controls, rate limits, and budget caps operational. Graceful degradation to lower-cost layers defined. Trust Cascade routing configured. What happens when the agent fails? The answer can't be "the network goes down."

"A network agent should not launch until all five gates are cleared. In Indian telecom, this isn't optional — it's the [minimum bar for AN Level 4 operations](#) and regulatory compliance."

Agent Governance Infrastructure for Indian Telecom

Rotavision provides the complete agent governance infrastructure for Indian telecommunications. Five products built from first principles for network agent operations, AN Level compliance, and billion-subscriber scale.

Orchestrate

Multi-Agent Network Operations and AN Governance

Enterprise-grade agent orchestration with Trust Cascade routing, policy enforcement, and bounded autonomy for the NOC. Agent fleet management, intent-to-policy translation, reasoning capture, and closed-loop control. The operational backbone for governed AN Level 4 deployment.

Guardian

Agent Reliability and Drift Monitoring

Continuous production monitoring for network agent behaviour. Catches drift, hallucination, and twin-production divergence before they impact operations. Identical monitoring in twin and production environments. 96% detection accuracy at less than 50ms overhead.

AgentOps

Enterprise Agent Registry and Policy Engine

From RotaScale. Centralised agent registry with identity, autonomy levels, blast radius limits, and risk classification. Declarative policy engine enforced at runtime. Flight recorder for every agent decision. The control plane for enterprise network agent operations.

Vishwas

Agent Fairness for Customer Decisions

Service quality fairness monitoring across circles, languages, and subscriber segments. Detects when agents systematically underserve rural areas or flag regional language users as fraud. Customer-facing explainability in 22 languages. TRAI QoS compliance validation.

Sankalp

Sovereign AI Gateway

Route agent traffic through India-hosted infrastructure with TRAI compliance built in. Automatic AI disclosure injection. Consent Registry integration. Data never leaves India. Strategic national infrastructure demands sovereign AI deployment.

BUILT FOR INDIAN TELECOM. AGENT-FIRST.

Your infrastructure. On-premise, private cloud, or hybrid. **No data leaves India.** Every product built for agent governance in telecom at billion-subscriber scale. AN Level 4 compliant from day one.

What Telcos Build With Agent Governance

Production agent systems making network decisions across fault prediction, capacity optimisation, fraud detection, and customer service. Each implementation demonstrates what becomes possible when network agents have proper operations infrastructure.

Autonomous Fault Prediction and Self-Healing

Network agents predict faults across radio, transport, and core domains. Self-healing workflows trigger automated remediation within blast radius limits. Guardian monitors every agent for drift. Orchestrate enforces policy boundaries. Full reasoning capture for every autonomous action.

Result: 45% reduction in mean-time-to-repair with zero unconstrained cascades

5G Slice Management with Governed Agents

Agents manage slice allocation in real time — dynamically provisioning capacity for enterprise, consumer, and IoT slices. Bounded autonomy prevents over-provisioning. Trust Cascade routes routine allocation through rules, complex decisions through agents. Cost awareness at every layer.

Result: 30% improvement in spectrum efficiency with auditable decisions

Multi-Agent Fraud Detection

Specialised agents — SIM swap pattern, CLI verification, call-flow analysis — each reason independently, then a coordinator synthesises signals. Multi-language spam detection catches fraud in Hindi, Tamil, Bengali, and regional dialects. DoT integration for cross-operator coordination. Complete audit trails.

Result: 78% fraud detection rate with 49% fewer subscriber complaints

Vernacular Customer Service Agents

Customer-facing agents handle plan queries, billing disputes, and number portability in 22 languages. TRAI-compliant AI disclosure and consent. Trust Cascade routes 75% to IVR, 15% to ML, 10% to agents. Vishwas monitors every recommendation for fairness across circles and demographics.

Result: 60% cost reduction with higher CSAT in regional languages

Digital Twin Governance Parity

Same policy engine and reasoning capture for agents in twin and production environments. Integrates with JioBrain, HCL ANA, or existing digital twin platforms. Validation in the twin matches production reality. Twin-production divergence monitoring prevents false confidence.

Result: Zero production surprises from twin-validated agents

AN Level Compliance Dashboard

Real-time agent governance monitoring mapped to TM Forum AN assessment framework. Agent maturity scoring per AN Level. TRAI QoS compliance tracking. Board-ready autonomous network governance reports. Agent-level drill-down with full reasoning traces.

Result: AN Level assessment preparation reduced from months to weeks

"The platform doesn't replace your network AI strategy — it makes your agents **production-ready for AN Level 4**. Same capabilities, but with the governance infrastructure regulators and the TM Forum expect."

AN Governance Accelerator

A combined assessment, platform, and integration package that maps your agent governance maturity to the TM Forum AN framework — and builds the governance layer to reach Level 4.

What's Included

1 AN Level Assessment

Agent governance maturity audit mapped to TM Forum AN Levels. Gap analysis with actionable roadmap from your current state to Level 4. Board-ready presentation with prioritised remediation plan and investment sizing.

2 Agent Registry and Policy Engine

Orchestrate + AgentOps configured for NOC workflows. Agent registration with autonomy levels, blast radius controls, and policy enforcement for network agents. Intent-to-policy translation for business-level guardrails.

3 Digital Twin Governance Parity

Same policy engine and reasoning capture for agents in twin and production environments. Integrates with JioBrain, HCL ANA, or your existing digital twin platform. What you validate in the twin is what you deploy in production.

4 Intent-to-Policy Translation

Convert business intent ("prioritise video in Mumbai during IPL") into agent policy guardrails. The bridge between what you want the network to do and what agents are allowed to do. Declarative, version-controlled, auditable.

5 OSS/BSS Integration

Pre-built connectors for telecom operations systems. Agent governance layer that sits alongside your existing NOC stack, not instead of it. Integration with network management, service assurance, and billing platforms.

Platform Stack

Agent orchestration Orchestrate	Reliability monitoring Guardian
Agent registry and policy AgentOps (RotaScale)	Sovereign gateway Sankalp

Engagement Options

<p>ASSESSMENT</p> <p>Rs 15L</p> <p>2 weeks. AN Level gap analysis. Agent governance audit. TM Forum alignment. Board-ready roadmap.</p>	<p>ACCELERATOR</p> <p>Rs 30L</p> <p>4 weeks. Full AN governance alignment. Agent registry setup. Digital twin parity. Executive presentation.</p>	<p>PRODUCTION</p> <p>Rs 50L+</p> <p>10-20 weeks. Full platform deployment. OSS/BSS integration. NOC training. Go-live support.</p>
--	--	---

Level 4 autonomy means agents decide without humans. The question is whether you'll get there with governance.

1.15 billion mobile connections. Network agents making decisions at machine speed. Jio and Airtel racing to AN Level 4. The agents are already deployed. The operations layer is what's missing. Without agent governance, Level 4 isn't autonomy — it's liability.

We'd like to show you where you stand. A 30-minute AN Level assessment — not a sales pitch — to benchmark your agent governance maturity against the TM Forum framework and identify your highest-value opportunities for governed autonomy.

[Request AN 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

Delhi | Bengaluru | Mumbai