

TELECOMMUNICATIONS INDUSTRY SOLUTION

The Intelligent Network

*From network digital twins to autonomous operations.
A strategic guide to AI-driven network intelligence for
the 5G/6G era.*

EXECUTIVE SUMMARY

Networks are becoming too complex for humans to manage alone. Graph Neural Networks now model network topology. Digital twins simulate changes before deployment. Self-healing networks detect and remediate issues autonomously. The path to Autonomous Network Operations (ANO) isn't science fiction - it's engineering. This guide shows how leading carriers are building the AI infrastructure for network autonomy.

WHAT'S INSIDE

- 01 The Complexity Crisis: Why networks need AI
- 02 Network Digital Twins: Graph models emerge
- 03 Self-Healing Networks: Autonomous remediation
- 04 Autonomous RAN: Intent-based management
- 05 Change Simulation: Test before deploy
- 06 Root Cause Analysis: AI-powered troubleshooting
- 07 Network Slicing: AI for 5G service assurance
- 08 The Platform: Rotascale for telecommunications
- 09 Getting Started: Engagement options

73%

AI QUALITY ISSUES REPORTED

40%

MTTR REDUCTION POSSIBLE

L4

ANO AUTONOMY TARGET

2030

6G COMMERCIAL EXPECTED

Why Networks Need AI Now

Modern telecommunications networks have crossed a complexity threshold. Millions of parameters. Billions of connections. Real-time optimization requirements that humans simply cannot meet manually.

<p>5G Complexity Explosion</p> <p>Network slicing, massive MIMO, dynamic spectrum sharing. A single cell site has hundreds of optimization parameters.</p>	<p>Real-Time Requirements</p> <p>5G latency SLAs in single-digit milliseconds. Optimization decisions must happen at machine speed.</p>
<p>Multi-Domain Integration</p> <p>RAN affects core. Core affects transport. Coordinated optimization requires unified AI-driven visibility.</p>	<p>Service Differentiation</p> <p>Network slicing promises different SLAs on shared infrastructure. eMBB, URLLC, mMTC.</p>

The ANO Maturity Journey

TM Forum defines five levels of Autonomous Network Operations. Most carriers are at L1-2. L4-5 requires AI infrastructure that doesn't exist today.

Level	Description	Human Role	AI Capability
L1 - Assisted	AI assists human operators	Executes all actions	Recommendations only
L2 - Partial	AI executes routine tasks	Monitors, approves complex	Rule-based automation
L3 - Conditional	AI handles most scenarios	Handles exceptions	ML-based decisions
L4 - High	AI manages complex situations	Strategic oversight only	Multi-agent reasoning
L5 - Full	Full autonomous operations	Business governance	Self-optimizing systems

ROTASCALE PLATFORM **Infrastructure for every ANO level**

Orchestrate provides multi-agent coordination for L3+. Guardian monitors AI decision quality. Steer enforces safety boundaries.

Orchestrate
Guardian
Steer
AgentOps

Graph Models for Network Intelligence

Networks are graphs. Nodes, edges, relationships. Graph Neural Networks (GNNs) can model network topology natively, learning patterns that traditional ML approaches miss. This is the foundation of the network digital twin.

<p>94%</p> <p>TOPOLOGY CHANGE PREDICTION</p>	<p>10x</p> <p>FASTER SIMULATION</p>	<p>1M+</p> <p>NODES MODELED</p>
---	--	--

Traditional ML treats network data as flat tables. GNNs understand topology—they learn that a router failure affects downstream nodes and that congestion propagates.

<p>Topology Awareness</p> <p>GNNs encode network structure directly via message passing. Learned embeddings capture connectivity patterns and propagation dynamics.</p>	<p>Multi-Hop Reasoning</p> <p>A congested backhaul affects dozens of downstream cells. GNNs aggregate information across graph neighborhoods, capturing cascading effects.</p>
<p>Dynamic Adaptation</p> <p>Networks change constantly. GNNs can be updated incrementally as topology evolves, without full retraining.</p>	<p>Explainable Predictions</p> <p>GNN attention weights show which paths and nodes contributed to predictions. Explainability critical for operator trust.</p>

Digital Twin Architecture

- ✓ **Topology Layer:** Graph representation of physical and logical network
- ✓ **State Layer:** Real-time telemetry mapped to graph nodes and edges
- ✓ **Model Layer:** GNN models for prediction, anomaly detection, optimization
- ✓ **Simulation Layer:** What-if analysis and change impact prediction

ROTASCALE PLATFORM

Data foundation for network digital twins

Context Engine ingests topology and builds graph-native representations. Guardian monitors GNN accuracy. Eval validates predictions against real outcomes.

Context Engine
Guardian
Eval

Autonomous Detection and Remediation

Self-healing networks detect issues, diagnose root causes, and implement fixes - all without human intervention. This isn't about replacing network engineers. It's about handling the volume of issues that humans can't address in real-time.

<p>Detect</p> <p>Anomaly detection across millions of metrics. Pattern recognition that identifies issues before they become outages. Correlation across domains.</p>	<p>Diagnose</p> <p>Root cause analysis using GNN-based reasoning. Trace fault propagation through network graph. Identify actual cause vs. symptoms.</p>	<p>Remediate</p> <p>Automated fix deployment within policy guardrails. Traffic rerouting. Parameter adjustment. Failover activation. Human escalation when needed.</p>
--	---	---

The Self-Healing Loop

- 1 Anomaly Detection**

ML models continuously analyze telemetry. Statistical anomalies, pattern deviations, threshold breaches detected in real-time. Digital twin provides context.
- 2 Impact Assessment**

GNN-based propagation modeling. Which services affected? Which customers impacted? What's the blast radius? Severity classification drives response priority.
- 3 Root Cause Isolation**

Multi-agent reasoning correlates symptoms across domains. Graph attention identifies fault origin. Eliminates false positives from correlated symptoms.
- 4 Remediation Selection**

Policy-constrained action selection. What can be done automatically? What needs approval? Simulate fix impact before deployment using digital twin.
- 5 Execution & Verification**

Deploy fix through orchestration layer. Monitor impact. Verify resolution. Rollback if metrics don't improve. Document for learning.

"The goal isn't zero human involvement. It's the right human involvement - strategic decisions, not routine firefighting."

ROTASCALE PLATFORM **Closed-loop healing with safety guardrails**

Orchestrate coordinates detection, diagnosis, and remediation agents in a safe execution pipeline. Steer enforces policy boundaries on what autonomous actions are permitted. Guardian monitors remediation effectiveness and triggers rollback when fixes don't improve metrics. AgentOps captures every healing action for post-incident analysis.

Orchestrate Steer Guardian AgentOps

Intent-Based Radio Management

The Radio Access Network is where complexity hits hardest. Thousands of parameters per cell. Interference coordination. Load balancing. Handover optimization. Intent-based networking lets operators specify "what" without prescribing "how."

From Commands to Intents

Traditional	Intent-Based
Set cell 1234 power to 43 dBm	"Optimize coverage in downtown"
Configure X2 handover to cell 5678	"Minimize dropped calls on Highway 101"
Enable carrier aggregation bands 1,3,7	"Maximize throughput for enterprise campus"
Set SINR threshold to -6 dB	"Maintain 99.9% reliability for URLLC slice"

"Intent-based networking separates business requirements from technical implementation. AI bridges the gap."

The RAN Autonomy Stack

Intent Engine

Translates business intents into technical objectives. Natural language understanding. Conflict resolution.

Policy Manager

Guardrails on what AI can change. Regulatory constraints. Vendor limitations. Safety boundaries.

Optimization Engine

Reinforcement learning for parameter tuning. Multi-objective optimization. Real-time adaptation.

Execution Layer

Safe deployment of changes. Rollback capability. Gradual rollout. Impact verification.

Key Use Cases

- ✓ Coverage and capacity optimization
- ✓ Energy efficiency (lower power during low traffic)
- ✓ Interference management in dense deployments
- ✓ Mobility optimization (handover thresholds)
- ✓ Dynamic spectrum sharing

ROTASCALE PLATFORM

Intent translation with safety boundaries

Orchestrate translates business intents into coordinated agent actions across the RAN. Steer enforces vendor constraints, regulatory limits, and safety boundaries on parameter changes. Guardian monitors optimization outcomes and detects when AI decisions degrade network performance.

Orchestrate

Steer

Guardian

Test Before You Deploy

Network changes cause outages. Configuration errors. Unexpected interactions. Capacity miscalculations. AI-powered change simulation lets you test configurations against your digital twin before touching production.

<p>70%</p> <p>OUTAGES FROM CONFIG CHANGES</p>	<p>85%</p> <p>PREVENTABLE WITH SIMULATION</p>	<p>3x</p> <p>FASTER CHANGE VELOCITY</p>
--	--	--

Simulation Workflow

Propose → Sync Digital Twin → Simulate Impact (GNN) → Deploy or Block

What Gets Simulated

- ✓ **Traffic Impact:** Congestion points, latency, traffic redistribution
- ✓ **Failure Scenarios:** Cascading failure analysis, redundancy validation
- ✓ **Service Impact:** SLA prediction, revenue risk quantification

AI Change Agents

- ✓ **Config Validator:** Syntax, policy compliance, vendor constraints
- ✓ **Impact Analyzer:** GNN-based topology effect prediction
- ✓ **Risk Assessor:** Probability/severity, change window recommendations
- ✓ **Rollback Planner:** Pre-computed recovery plans

ROTASCALE PLATFORM **Simulate with confidence, deploy with safety**

Context Engine maintains the digital twin. Orchestrate coordinates change agents. Steer enforces policies. Guardian validates outcomes.

Context Engine
Orchestrate
Steer
Guardian

AI-Powered Troubleshooting

When something breaks, finding the root cause fast is critical. Traditional troubleshooting is slow - checking logs, correlating events, testing hypotheses. AI agents can reason about failures across the entire network graph in seconds.

<p>40%</p> <p>MTTR REDUCTION</p>	<p>90%</p> <p>FIRST-CALL RESOLUTION</p>	<p>10K+</p> <p>EVENTS CORRELATED/SEC</p>
---	--	---

Multi-Agent Root Cause Analysis

Complex failures require multiple perspectives. Our RCA system uses coordinated agents, each specialized for different aspects of failure analysis.

<p>Symptom Collector</p> <p>Gather All Evidence</p> <p>Aggregates alarms, metrics anomalies, customer complaints, and automated tests. Normalizes data from different sources into common format. Establishes timeline.</p>	<p>Correlation Engine</p> <p>Connect the Dots</p> <p>Uses GNN to trace symptom relationships through network topology. Identifies common ancestors. Separates causes from effects. Eliminates noise from coincidental events.</p>
<p>Hypothesis Generator</p> <p>What Could Cause This?</p> <p>Generates ranked list of potential root causes. Considers recent changes. Checks known issue patterns. Incorporates historical failure modes.</p>	<p>Validator Agent</p> <p>Test Each Hypothesis</p> <p>Runs targeted diagnostics to confirm or eliminate hypotheses. Queries affected components. Checks configurations. Validates connectivity.</p>

RCA Output

- ✓ **Root Cause:** Specific component or configuration issue
- ✓ **Evidence Chain:** How symptoms trace back to cause
- ✓ **Impact Assessment:** Services and customers affected
- ✓ **Remediation Options:** Ranked by speed and safety
- ✓ **Prevention:** Recommendations to avoid recurrence

"The AI found the root cause in 90 seconds. It would have taken our team 2 hours. Not because they're slow - the network is just too complex."

— NOC Manager, Regional Carrier

ROTASCALE PLATFORM From alarm storms to root cause in seconds

Orchestrate coordinates symptom collector, correlation engine, hypothesis generator, and validator agents. Context Engine provides the GNN-based network graph for correlation. AgentOps captures the full RCA reasoning chain for review.

Orchestrate
Context Engine
AgentOps

AI for 5G Service Assurance

Network slicing promises differentiated services on shared infrastructure. But delivering different SLAs simultaneously requires real-time resource allocation that only AI can provide. This is where the rubber meets the road for 5G monetization.

<p>eMBB Slice</p> <p>Enhanced Mobile Broadband</p> <p>High throughput. Consumer video. Enterprise data. Best-effort with minimum guarantees. Peak rate optimization.</p>	<p>URLLC Slice</p> <p>Ultra-Reliable Low Latency</p> <p>Mission-critical. Industrial automation. Remote surgery. 99.999% reliability. <10ms latency. Zero tolerance for degradation.</p>	<p>mMTC Slice</p> <p>Massive Machine-Type</p> <p>IoT scale. Millions of devices. Low power. Small packets. Massive connection density. Efficiency critical.</p>
--	---	---

AI-Driven Slice Management

<p>Resource Allocation</p> <p>RL agents optimize spectrum, compute, and network resources across slices in real-time. Maintain SLAs while maximizing utilization.</p>	<h3>The Challenge</h3>																
<p>SLA Monitoring</p> <p>Continuous verification that each slice meets its contracted performance. Early warning when SLAs at risk. Proactive remediation.</p>	<p>Each slice has different requirements. Resources are shared. Optimization is multi-objective with hard constraints.</p>																
<p>Admission Control</p> <p>AI-driven decisions on slice instantiation. Can the network support this new slice? Impact on existing slices?</p>	<table border="1"> <thead> <tr> <th>Metric</th> <th>eMBB</th> <th>URLLC</th> <th>mMTC</th> </tr> </thead> <tbody> <tr> <td>Priority</td> <td>Throughput</td> <td>Latency</td> <td>Density</td> </tr> <tr> <td>Reliability</td> <td>99.9%</td> <td>99.999%</td> <td>99%</td> </tr> <tr> <td>Latency</td> <td><50ms</td> <td><1ms</td> <td><1s</td> </tr> </tbody> </table>	Metric	eMBB	URLLC	mMTC	Priority	Throughput	Latency	Density	Reliability	99.9%	99.999%	99%	Latency	<50ms	<1ms	<1s
Metric	eMBB	URLLC	mMTC														
Priority	Throughput	Latency	Density														
Reliability	99.9%	99.999%	99%														
Latency	<50ms	<1ms	<1s														
<p>Cross-Slice Optimization</p> <p>Global optimization across all slices. Resource sharing when possible. Isolation when required. Pareto-optimal allocation.</p>	<p><i>"Manual slice management doesn't work. The dynamics are too fast. The interactions too complex. AI is required."</i></p>																

ROTASCALE PLATFORM

SLA assurance across every slice

Guardian monitors SLA compliance per slice and alerts before violations. Orchestrate coordinates RL agents for cross-slice optimization. Steer enforces isolation policies. Accelerate optimizes inference latency for real-time slice management.

Guardian

Orchestrate

Steer

Accelerate

Rotascale for Telecommunications

The Trust Intelligence Platform provides the complete stack for production-grade telecom AI. From network digital twins to autonomous operations to agent governance.

Guardian

AI Reliability Monitoring

Continuous monitoring of AI model performance in network operations. Detect drift as network conditions change. Alert when models degrade.

- ✓ Model accuracy tracking by use case
- ✓ Prediction confidence calibration
- ✓ Drift detection and alerting
- ✓ Fairness monitoring across network segments

Orchestrate

Multi-Agent Platform

Coordinate network AI agents safely. RCA agents, optimization agents, simulation agents. Full context preservation across agent boundaries.

- ✓ Agent coordination patterns
- ✓ Safe multi-agent execution
- ✓ Human escalation triggers
- ✓ Complete reasoning capture

Steer

Runtime Behavior Control

Guardrails that enforce safety in network AI. Prevent dangerous reconfigurations. Enforce change windows. Limit blast radius.

- ✓ Network safety guardrails
- ✓ Change policy enforcement
- ✓ Blast radius limiting
- ✓ Rollback triggers

AgentOps

Enterprise Agent Governance

Complete visibility into every AI agent in your network operations. From optimization bots to diagnostic agents. Registry and audit trail.

- ✓ Universal agent registry
- ✓ Decision audit trails
- ✓ Policy compliance enforcement
- ✓ Cost attribution

Network Data Foundation

Network telemetry from millions of elements. Topology from multiple OSS systems. Our Data Intelligence capabilities unify it for AI consumption.

Graph Data Layer

Native graph storage for network topology. GNN-ready data representation. Real-time updates.

Telemetry Integration

Streaming ingestion from network elements. Time-series optimization. Multi-vendor support.

Context Engine

Unified view across OSS/BSS systems. Relationship-preserving data assembly. AI-ready context.

USE CASES

Telecom AI Applications

The Trust Intelligence Platform powers AI across telecom operations - from network optimization to customer experience to fraud prevention.

Network Digital Twin

GNN-based digital twin of your entire network. Topology-aware modeling. Change simulation. What-if analysis. Real-time synchronization with production.

Products: Orchestrate, Context Engine, Guardian

Self-Healing Networks

Autonomous detection, diagnosis, and remediation. AI agents that can identify root causes and implement fixes within policy guardrails.

Products: Orchestrate, Steer, AgentOps

RAN Optimization

Intent-based radio management. RL agents that continuously optimize coverage, capacity, and efficiency. Human oversight for strategic decisions.

Products: Orchestrate, Guardian, Steer

Network Slicing Assurance

AI-driven resource allocation across slices. SLA monitoring and enforcement. Dynamic optimization as demand changes.

Products: Guardian, Orchestrate, Context Engine

Customer Experience AI

AI-powered customer service for network issues. Intelligent troubleshooting. Proactive outreach. NPS improvement through faster resolution.

Products: Orchestrate, Guardian, Steer

Fraud Detection

Real-time detection of subscription fraud, SIM swap attacks, and revenue assurance issues. Graph-based analysis of calling patterns.

Products: Guardian, Context Engine

"We've been trying to implement AI in network operations for years. Rotascale is the first platform that actually works in production."

— CTO, Regional Carrier

LOOKING AHEAD

Preparing for 6G

6G is expected to reach commercial deployment around 2030. The complexity will be an order of magnitude greater than 5G. AI isn't optional for 6G - it's architectural. Building AI capabilities now is preparing for a future where networks manage themselves.

What 6G Brings

- TERAHERTZ SPECTRUM**

Beyond mmWave

New spectrum bands with extreme propagation challenges. Requires real-time beam management that only AI can provide.
- AI-NATIVE ARCHITECTURE**

AI Built In, Not Bolted On

6G standards incorporate AI natively. Autonomous operation expected. Intelligence distributed throughout the network.
- SENSING + COMMUNICATION**

Joint Radar-Communication

Networks that sense their environment. Integrated positioning. New data sources for AI to process and act upon.
- NON-TERRESTRIAL INTEGRATION**

LEO Satellites + UAVs

3D network topology. Moving nodes. Extreme complexity. Only AI can manage the coordination requirements.

Why Build AI Capability Now

<p>Learning Curve</p> <p>AI for network operations is hard. Building expertise takes years. Start now to be ready for 6G.</p>	<p>Data Foundation</p> <p>AI needs training data. Historical network data you collect now trains 6G-ready models.</p>
<p>Process Integration</p> <p>Integrating AI into operations requires organizational change. Culture shifts take time.</p>	<p>Competitive Position</p> <p>Carriers who master AI for 5G will lead 6G. Those who wait will struggle to catch up.</p>

"6G will be impossible to operate without AI. The question isn't whether to invest in AI capability - it's whether you start now or wait until it's too late."

ROTASCALE PLATFORM

Build 6G-ready AI infrastructure today

The Rotascale platform is designed for increasing network complexity. Orchestrate scales from simple automation to full multi-agent autonomy. Guardian's monitoring framework adapts as new AI models and network technologies emerge. The investment you make now in AI trust infrastructure directly transfers to 6G operations.

Orchestrate

Guardian

AgentOps

Steer

Your Network AI Journey

Whether you're building network digital twins, implementing self-healing capabilities, or governing AI across your operations, we have engagement options designed for telecommunications.

ANO ASSESSMENT	DIGITAL TWIN PILOT	ANO PLATFORM
\$40K	\$150K	\$500K+
3-4 weeks	8-12 weeks	4-6 months
<ul style="list-style-type: none"> ✓ Current AI capability audit ✓ ANO maturity assessment ✓ Use case prioritization ✓ Implementation roadmap 	<ul style="list-style-type: none"> ✓ Network segment digital twin ✓ GNN model training ✓ Change simulation capability ✓ RCA use case demonstration 	<ul style="list-style-type: none"> ✓ Full network digital twin ✓ Self-healing implementation ✓ Agent governance framework ✓ Ongoing advisory

What You Get

- ✓ **Network-native architecture** - Built for telecom from day one. Graph models, real-time telemetry, multi-domain integration.
- ✓ **Production-grade reliability** - AI that works in carrier networks. Proven at scale. Enterprise SLAs.
- ✓ **Vendor agnostic** - Works with your existing OSS/BSS. Multi-vendor network support. Standards-based integration.
- ✓ **6G ready** - Architecture designed for increasing complexity. Building blocks for autonomous operations.

Who We Work With

Network Leadership

CTOs, VPs of Network Engineering, and Directors of Network Operations building AI-driven network capabilities.

Technology Leaders

CIOs and VPs of IT implementing AI infrastructure that integrates with network operations.

Ready for autonomous network operations?

Start with an assessment. Understand your ANO maturity and the path to Level 4.

[Request Assessment](#)

[Learn More](#)

Contact

contact@rotascale.com · +1 (415) 524-0007
 rotascale.com/solutions/telecommunications