



AIGW Policy Observatory
Global AI Governance & Workforce Transformation

Global AI Governance and Workforce Transformation Policy Observatory

AIGW WORKFORCE TRANSITION DEEP DIVE

Deutsche Telekom: When AI Agents Start Acting on the Network

How agentic AI moved network operations from monitoring and recommendation toward autonomous corrective action — and why operational authority now matters as much as productivity

3 October 2026

A decorative graphic in the bottom right corner consisting of several concentric, curved lines in shades of blue and purple, creating a sense of motion or a network structure.

CONTENTS

Contents

Executive Summary and Key Findings	03
The RAN Guardian in Production	09
Authority and Agent Governance	14
Work and Workforce Transition	18
Autonomous Networks and Operating Context	20
Measured Evidence and Comparison	28
Evidence Limits and Research Agenda	32
Conclusion and Notes	36

EXECUTIVE SUMMARY

Executive Summary

Deutsche Telekom's RAN Guardian Agent is an important boundary case in the evolution of AI-enabled work.

The system does not merely summarize information for an engineer. It does not stop at detecting a problem. It can identify an operational situation, analyze network conditions, recommend a response and autonomously execute corrective actions on a live mobile network, including reallocating network resources and adjusting network configurations.[1]

That change matters because it moves AI from assistance into delegated operational action.

In November 2025, Deutsche Telekom announced that the RAN Guardian Agent had gone live in German mobile-network management. The system operates as a multi-agent architecture. One agent scans public information sources for events likely to create unusual network demand. Another assesses local network capacity, nearby antennas and real-time performance parameters. A further agent can execute optimization measures and document what happened so the system can improve future responses.[1]

Deutsche Telekom reported that selected network-management processes which previously took employees roughly an hour could be completed within a few minutes.[1]

Follow-on reporting in February 2026 made the production evidence stronger. Deutsche Telekom said that in the first month after launch the RAN Guardian autonomously triggered more than 100 remediation actions around Christmas-market events. The company reported that live operations had reduced the time required to manage major events from hours to around one minute — an improvement of more than 95% — and that the system had identified roughly 237,000 events for 2026. During the February Carnival season it identified around 130 events and parades with more than 10,000 expected participants, pre-checked 611 mobile sites, and adjusted the five sites that encountered peak loads.[2]

This is unusually strong evidence for agentic AI because the outcome is not merely:

- text generated;
- recommendation produced;
- employee time estimated.

The system is acting on operational infrastructure.

That makes the workforce question more complex.

The strongest documented work transition is:

manual monitoring, event identification, network analysis and corrective execution → multi-agent monitoring, analysis and autonomous remediation

EXECUTIVE SUMMARY

The public evidence does not show what happened to the number of network experts employed, whether roles were removed, whether employees were redeployed or whether specific reskilling programmes were created for RAN Guardian users.

The correct workforce classification is therefore not "AI replaced network engineers."

It is:

AI absorbed a meaningful share of operational execution and compressed the cycle time of network-management work, while the longer-term workforce response remains publicly under-specified.

The case is also important because it exposes the limits of the phrase human in the loop.

Deutsche Telekom says the RAN Guardian works through intelligent interaction between AI and network experts. But the November 2025 production announcement does not specify whether a human must approve every configuration change, whether network experts can veto or reverse actions, what thresholds trigger escalation, or how authority changes when an action is higher risk.[1]

That absence should not be filled with assumption.

The system's autonomy is directly evidenced.

Its complete human-control architecture is not.

This creates what AIGW calls the Authority Gap:

When autonomous action is publicly evidenced but the corresponding human approval, override and accountability mechanisms are not described at the same level of specificity.

Later Deutsche Telekom material shows that the company is thinking explicitly about this problem at enterprise level. In September 2026, Deutsche Telekom experts described a risk-based agent-governance model in which agents have identities, access rights, logged actions and boundaries; particularly critical actions can be configured to require human approval, and responsibility remains with the organization deploying the agent. [3][4]

That is highly relevant follow-on governance evidence.

But it should not be retroactively treated as proof that the RAN Guardian used all of those controls in November 2025.

The distinction is central to rigorous analysis:

enterprise governance philosophy ≠ use-case-specific decision-right evidence

The technology trajectory is also expanding.

EXECUTIVE SUMMARY

In February 2026 Deutsche Telekom and Google Cloud announced MINDR — Multi-Agentive Intelligent Network Diagnostics & Remediation — a broader system intended to move agentive operation from the radio access network into end-to-end service operations across RAN, transport and core network domains.[2]

MINDR is designed to move from:

reactive troubleshooting

toward:

predictive, service-oriented diagnosis and autonomous remediation before customers experience the problem.

Initial production releases were planned for later in 2026.[2][5]

The distinction between RAN Guardian and MINDR matters.

RAN Guardian provides production evidence.

MINDR provides evidence of the next operating-model direction.

Together they show a progression:

monitor → analyze → recommend → execute → coordinate across domains

As autonomy expands, the work of network experts is likely to move away from manually performing every operational step and toward activities such as:

- policy and intent definition; • exception management; • validation; • AI-agent supervision; • system assurance; • cross-domain diagnosis; • risk control.

However, those capability shifts remain analytical implications unless directly evidenced for Deutsche Telekom's network workforce.

The case therefore supports a broader proposition:

Agentive AI changes work most profoundly not when it generates a better recommendation, but when an organization delegates the authority to act.

Once that threshold is crossed, the relevant management question changes from:

"How accurate is the AI?"

to:

"What is the AI allowed to do, under which conditions, who can stop it, and who remains accountable?"

KEY FINDINGS

Key Findings

01

1. RAN Guardian is a production agent, not a conventional copilot

Deutsche Telekom says the system is live and can autonomously implement corrective network actions.
[1]

This distinguishes the case from AI systems limited to recommendation or first-pass text production.

02

2. The workflow has moved from sequential human analysis toward closed-loop agentic operation

The system identifies events, evaluates network load, monitors parameters, recommends measures, executes changes and records outcomes.[1]

The transition is therefore larger than task automation.

It is a redesign of the operating loop.

03

3. Process compression is substantial

Deutsche Telekom initially reported selected processes falling from about an hour to a few minutes.[1]

By February 2026, the company reported management of major events falling from hours to around a minute, or more than 95%.[2]

These are company-reported operational outcomes, not independent standardized benchmarks.

04

4. Autonomous action is directly evidenced

The agent can reallocate network resources and adjust configurations on the live network.[1][2]

That makes decision rights a first-order governance issue.

KEY FINDINGS

01

5. Human interaction is described, but use-case-specific authority remains incomplete

The company refers to interaction between network experts and AI.[1]

The production source does not publish a complete matrix of:

- approval; • veto; • override; • rollback; • escalation; • final accountability.

03

7. Workforce consequences should not be inferred from process automation

No reviewed RAN Guardian source directly documents:

- layoffs; • reduced hiring; • redeployment; • reskilling; • changed network-operations headcount.

The workforce response remains under-evidenced.

02

6. The next stage expands autonomy across network domains

MINDR extends the agentic model from RAN toward RAN + transport + core and is designed for predictive, end-to-end service operations.[2][5]

This increases both the potential value and the governance surface.

04

8. Deutsche Telekom has a broad AI-skills infrastructure, but it should not be silently attributed to this network-operations episode

In 2025, Deutsche Telekom reported 30,000 employees trained in AI-related topics, 52,000 participations in AI-related Learning from Experts sessions, and a goal for 90% of employees outside T-Mobile US to use AI tools regularly in 2026.[6]

This establishes enterprise capability-building context.

It does not establish a RAN Guardian-specific reskilling programme.

CASE CONTEXT

Why this case matters

01

Why This Case Matters

For most of the recent history of enterprise AI, the dominant interface was advisory.

AI could:

- predict; • classify; • summarize; • recommend; • prepare a first version.

A human remained the actor.

Agentic systems change that architecture.

They can connect perception, reasoning and action.

In a network environment, that distinction is operationally important.

A system that says:

"This cell is likely to become congested"

is not equivalent to a system that:

reallocates resources or changes the network configuration itself.

The first changes analysis.

The second changes authority.

Telecommunications also makes a useful test environment because networks are:

- large; • dynamic; • multi-domain; • real-time; • highly interdependent; • difficult to manage manually at full scale.

Network operations therefore expose both the promise and the risk of agentic AI.

The promise is faster:

- detection; • correlation; • diagnosis; • remediation.

The risk is that an erroneous action can itself affect service.

This means operational AI cannot be governed only through model accuracy.

It also requires control over:

- identity; • permissions; • action scope; • logging; • validation; • escalation; • rollback; • accountability.

RAN Guardian is valuable because it shows the technology crossing that threshold in production.

02

The Starting Point: Network Operations as Context-Heavy Human Work

Traditional network operations already contain substantial automation.

Telecom operators have long used:

- alarms; • dashboards; • threshold rules; • orchestration systems; • scripts; • automated provisioning; • performance-management tools.

What remains difficult is not simply collecting telemetry.

The difficult work is often context construction.

A network expert may need to understand:

- whether an alarm represents a real service risk; • whether several alarms share one root cause; • whether a traffic spike is temporary or persistent; • whether a nearby site can absorb additional load; • whether a configuration change will solve one problem while creating another; • whether the situation is normal for a concert, stadium event or commuter pattern; • whether action is needed now or can wait.

TM Forum's 2026 work on context management for AI-native operations describes this problem directly. It notes that much operational context is still created and interpreted by humans: network-operations-center engineers investigate alarms, correlate events and add missing information based on experience. As networks become more automated, the industry is trying to make that context explicit, governed and usable by both people and AI agents.[7]

This makes RAN Guardian especially interesting.

The agent is not only automating a known manual button-click.

It is automating parts of the context-building process that previously made expert action possible.

PRODUCTION ARCHITECTURE

The RAN Guardian Architecture: Three Agentic Functions

Deutsche Telekom's public description allows the workflow to be reconstructed at a relatively high level.[1]

01

Stage 1 — Event discovery and classification

One agent continuously scans publicly available sources, including online directories and social-media sources, to identify events across Germany.

It:

- compiles events; • estimates scale; • determines location; • classifies them; • automatically verifies the gathered information.[1]

This replaces part of a planning function that otherwise depends on people knowing:

- what is happening; • where; • when; • how many people may attend.

The output is operational context.

02

Stage 2 — Capacity and anomaly analysis

A second component evaluates whether the local mobile network can handle expected traffic.

It considers:

- capacity of nearby antennas; • current network parameters; • utilization; • performance data; • inventory data; • coverage data.[1]

The system can recognize high utilization and propose optimization measures.

This compresses the diagnostic layer of network management.

03

Stage 3 — Corrective action

A further agent can execute the optimization.

Examples published by Deutsche Telekom include:

- reallocating network resources; • adjusting network configuration.[1]

The system also documents actions and outcomes so that future operation can improve.[1]

04

The important transition

The agentic chain is therefore:

discover context → assess network → identify risk → recommend action → execute action → record outcome

This is qualitatively different from:

dashboard → engineer notices issue → engineer investigates → engineer decides → engineer implements.

The human is no longer necessarily the executor of each step.

WORKFLOW REDESIGN

Before → Agentic Intervention → Emerging Workflow

FIGURE 1.1

AIGW workflow reconstruction

How network operations are changing

Traditional / more manual workflow	Agentic intervention	Emerging workflow
Staff identify major events relevant to network demand	Agent scans public sources and classifies events	Experts receive machine-generated event context at much larger scale ^[1]
Staff inspect capacity and telemetry across systems	Agent correlates network monitoring, inventory, performance and coverage data	AI performs first-pass situation assessment ^[1]
Staff investigate congestion risk	Agent predicts/detects anomalies and high utilization	Experts can focus on unusual or strategic cases ^[1]
Staff determine an optimization action	Agent recommends corrective measures	Recommendation generation is delegated ^[1]
Staff implement configuration/resource changes	Agent can autonomously execute changes	Human execution is no longer required for every action ^[1]
Staff record outcome and learn from incident	System documents actions and outcomes	Operational history can feed future agent decisions ^[1]

Source: AIGW reconstruction from Deutsche Telekom's public RAN Guardian description [1]

Note: Directional reconstruction; not a complete time-and-motion study.

PRODUCTION EVIDENCE

Production Evidence: From “One Hour” to Autonomous Remediation

The first production announcement in November 2025 reported that processes previously taking employees about an hour could be completed within a few minutes.[1]

At that point, the exact denominator and evaluation method were not public.

The statement is best treated as a representative workflow result.

01

Follow-on evidence strengthened the case

In February 2026, Deutsche Telekom provided more operational detail.[2]

It said that since the November launch:

- more than 100 remediation actions had been autonomously triggered during Christmas-market events in the first month;
- event-management time had fallen from hours to around a minute;
- the company characterized the improvement as more than 95%;
- the system had identified approximately 237,000 events for 2026;
- around 130 Carnival events and parades with more than 10,000 expected participants were identified;
- 611 mobile sites serving those events were pre-checked;
- five sites experienced peak loads and were optimized by RAN Guardian.[2]

A January 2026 network update separately described roughly 240,000 events identified/classified for 2026 and said this represented around 100 times as many events as could previously be identified and adjusted in the relevant process.[8]

These numbers are close but not identical.

AIGW should not manufacture a false reconciliation.

The correct approach is to preserve the source definitions:

- February 2026 reporting: ~237,000 events;
- January 2026 reporting: ~240,000 events;
- later festival reporting: ~1,750 open-air events benefiting from automated functions in a narrower festival/open-air context.[9]

These are not necessarily the same event population.

02

Why the “100 autonomous actions” metric matters

Many enterprise AI cases prove that a model can generate output.

This metric proves something different:

the production system was allowed to trigger operational remediation repeatedly.

That is the strongest evidence in the case.

OPERATIONAL AUTHORITY

What “Autonomous” Means Here

Autonomy is often used loosely.

In the RAN Guardian case, the public description supports at least three forms of autonomy.

01

Autonomous monitoring

The system continuously evaluates relevant information and network parameters.[1]

02

Autonomous analysis

It identifies anomalies, classifies problems and evaluates potential network-load risk.[1]

03

Autonomous execution

It can implement corrective actions such as resource reallocation or configuration changes.[1]

That final layer is decisive.

A system that monitors and recommends is still primarily an analytical tool.

A system that executes creates a machine actor inside the operating process.

04

AIGW distinction

AI Assistance Human acts; AI advises.

Agentic Delegation AI is authorized to perform a defined operational action.

RAN Guardian belongs in the second category.

05

What the workflow reconstruction does not establish

The public evidence does not provide a complete time-and-motion study of network experts before and after RAN Guardian.

It does support a directional reconstruction.

This table is an AIGW reconstruction based on Deutsche Telekom's public description.

It does not establish:

- the percentage of total network-operations work automated;
- how often humans intervene;
- how many actions require pre-approval;
- whether every network-management process follows this architecture.

NETWORK ROADMAP

From RAN Guardian to MINDR: Autonomy Expands Across Domains

RAN Guardian focuses on the Radio Access Network.

The next stage is broader.

In February 2026 Deutsche Telekom and Google Cloud announced MINDR — Multi-Agentive Intelligent Network Diagnostics & Remediation.[2]

The intended architecture extends agentic operations across:

- radio access network; • transport network; • core network.

The strategic shift is from domain-specific optimization toward service-oriented end-to-end operations.

01

Reactive → predictive

Traditional troubleshooting often begins after:

- an alarm; • service degradation; • a customer complaint.

MINDR is designed to:

- correlate signals across network domains; • build an end-to-end view of service performance; • identify service-impacting issues proactively; • analyze root causes; • support controlled remediation before customer impact.[2]

02

Multi-agent coordination

Deutsche Telekom says MINDR operates as a collaborative multi-agent system and intends to use Agent-to-Agent protocol coordination subject to availability.[2]

This matters because governance complexity increases when: one agent does not merely act, but several agents coordinate actions across different technical domains.

03

Production boundary

The February and March 2026 sources say initial MINDR deployments were planned for later in 2026.[2][5]

As of the public sources reviewed for this report, RAN Guardian has clear live-production evidence. MINDR should be described as the follow-on implementation direction, not silently treated as having the same mature production evidence.

HUMAN AUTHORITY

The Authority Gap

The system's autonomy is directly evidenced. Its complete use-case-specific human-control architecture is not.

RAN Guardian's public documentation tells us clearly what the AI can do.

It tells us less about what a person must authorize before the AI does it.

That asymmetry is analytically important.

“Human in the loop” tells us almost none of this.

For operational agents, the relevant question is:

Which authority layers remain human, and which are delegated?

EVIDENCE BOUNDARY

What the production source specifies

01

Directly evidenced

The system can:

- identify events; • analyze network behavior; • detect anomalies; • recommend measures; • autonomously implement corrective actions; • document actions and outcomes.[1]

02

Also evidenced

Deutsche Telekom describes an intelligent interaction between network experts and AI.[1]

03

Not specified in the production source

The source does not provide a formal answer to questions such as:

- Does every configuration change have a risk classification? • Which changes can occur without human approval? • Which changes require approval? • Can an engineer veto a planned action? • Can an engineer immediately override an action already executed? • Is rollback automatic or human-triggered? • What confidence threshold stops autonomous action? • What happens when agent recommendations conflict? • Who is operationally accountable for an erroneous autonomous change?

This is the Authority Gap.

It is not evidence that controls do not exist.

It is evidence that the public use-case documentation does not specify them.

04

The Human Authority Stack

The Kaiser Permanente case showed an AI system where human authority is relatively explicit.

Deutsche Telekom provides a useful contrast because autonomous execution is explicit while some use-case-specific authority details are not.

AIGW therefore proposes a Human Authority Stack for agentic operations.

GOVERNANCE CONTEXT

Deutsche Telekom's Later Agent-Governance Model

By September 2026, Deutsche Telekom was publicly describing a much more explicit governance architecture for AI agents in enterprise settings.[3][4]

Its experts argued that when many agents operate across an organization, each agent needs:

- a clear identity; • defined access rights; • logged actions; • boundaries it cannot cross.[3]

The company also described a risk-based control approach:

- agents can be tested by people before release; • security and compliance controls can be integrated; • actions can be classified by risk; • the system can determine when human intervention is required; • particularly critical actions can be configured to always require human approval; • cost and usage limits can be imposed.[3]

The most important governance statement is:

responsibility cannot be delegated to the AI agent.[3]

The organization deploying the agent remains responsible for deciding:

- which tasks may be automated; • which controls are required.

01

Why this is important — and why it must remain separate

This material demonstrates that Deutsche Telekom's broader agent strategy is moving toward risk-tiered autonomy.

That is strong institutional context.

But the enterprise T-AI Agentic Hub is not the same system as RAN Guardian.

Therefore the report should say:

Deutsche Telekom later articulated an agent-governance model that includes risk-based human approval and access controls.

It should not say:

RAN Guardian uses the T-AI Agentic Hub approval model.

That connection is not publicly established in the reviewed sources.

AIGW ANALYTICAL FRAMEWORK

The Human Authority Stack for Agentic Operations

FIGURE 2.1

AIGW framework

Seven questions that locate human authority

Authority layer	Operational question
Layer 1 — Scope permission	What actions is the agent authorized to perform at all? Example: • monitor only; • recommend; • adjust selected configuration variables; • execute broader remediation. ^[1]
Layer 2 — Trigger conditions	What conditions permit autonomous action? Examples: • low-risk routine; • known event type; • confidence above threshold; • no conflicting signals. ^[1]
Layer 3 — Pre-action approval	Does a human need to approve before execution? This can vary by risk. ^[1]
Layer 4 — Runtime monitoring	Can humans see: • what the agent is doing; • why; • on which systems; • with which permissions? ^[1]
Layer 5 — Override / stop	Can a human: • cancel; • pause; • reverse; • roll back? ^[1]
Layer 6 — Post-action verification	Who checks whether the autonomous action produced the intended network outcome? ^[1]
Layer 7 — Accountability	Which person or organizational role remains accountable? ^[1]

Source: AIGW analytical framework; the RAN Guardian source establishes autonomous action but not the full control matrix ^[1]

Note: This is an analytical framework, not a description of Deutsche Telekom's published RAN Guardian configuration.

WORKFORCE TRANSITION

Network work and capability building

01

Network Work Is Becoming an Intent-and-Exception Profession

The public evidence does not directly document a new Deutsche Telekom network-expert job description.

Any capability conclusion therefore needs to be framed as an implication.

But the workflow change points toward a plausible redistribution of work.

Work that the agent can absorb:

- continuous monitoring; • event discovery; • first-pass correlation; • utilization analysis; • known optimization actions; • action logging.

Intent definition:

What outcome should the network optimize for?

- quality; • resilience; • energy; • cost; • customer priority; • emergency service.

Policy design:

Which actions are allowed?

Under which limits?

Exception management:

What happens when:

- the event is novel; • signals conflict; • the model is uncertain; • the action has wider blast radius?

Assurance:

Did the agent's action actually improve service?

Did it create an unintended effect elsewhere?

System supervision:

Are multiple agents behaving coherently?

Are access rights still appropriate?

Strategic network engineering:

What architecture changes reduce the need for repeated remediation in the first place?

02

Enterprise-wide capability context

Deutsche Telekom's broader workforce strategy provides important context for the transition.

Its 2025 employee-development reporting says:[6]

- 30,000 employees were trained in AI-relevant topics in 2025; • 52,000 participations were recorded in AI-related Learning from Experts sessions; • 47,000 active users regularly used the Percipio learning platform; • the company aimed for 90% of employees outside T-Mobile US to use AI tools regularly in 2026; • the company aimed for 100% of those employees to possess necessary AI competence by 2028; • around 70,000 employees had completed the company's skill-management process; • the AI-based Growth Hub had 110,000 registered users at launch in November 2025.[6]

The company also reported that the share of digital experts outside T-Mobile US had risen to 22.7% in 2024 and targeted 25–30% by 2027.[6]

This demonstrates that Deutsche Telekom is not deploying AI in an organizational vacuum.

It has a scaled learning and skills infrastructure.

WORKFORCE EVIDENCE

Capability limits and company-wide workforce context

01

Enterprise Skills Context: Capability Building Exists, but the Network-Specific Link Is Missing

What this evidence supports:

It supports the conclusion that Deutsche Telekom is building:

- general AI fluency; • digital expertise; • skills visibility; • internal mobility infrastructure.

What it does not support:

It does not establish:

- how many network experts were trained specifically for RAN Guardian; • whether RAN Guardian changed a job family; • whether network operations employees were redeployed; • whether AI orchestration is formally part of their role profile; • whether certification is required to supervise autonomous agents.

AIGW research rule:

Do not use enterprise-wide AI training to fill a missing use-case-specific workforce response.

The correct formulation is:

Deutsche Telekom has broad organizational capability infrastructure that could support agentic work transitions, but the public RAN Guardian evidence does not disclose the specific workforce pathway.

02

Company-Wide Restructuring Is Real — and Must Be Kept Separate

Deutsche Telekom's workforce is also changing for reasons far broader than RAN Guardian.

The 2025 HR Factbook reports Group headcount of around 198,079 FTEs, broadly stable year on year.[10]

The Germany operating segment fell from 57,303 FTEs in 2024 to 55,089 in 2025, a decline of 3.9%.[10]

Across Germany more broadly, Deutsche Telekom reported 70,751 employees in 2025 compared with 74,550 in 2024.[11]

The company describes ongoing structural change and says it uses socially responsible restructuring instruments including:

- partial retirement; • committed retirement; • severance payments; • transfer management; • qualification support.[6][10]

Since 2018, Deutsche Telekom says around 3,400 employees in Germany have received career-change advice through transfer management and around 65% opened new career options through that advice.[6]

These are important workforce facts.

They should not be connected causally to RAN Guardian without evidence.

The company itself attributes segment-level headcount changes to multiple factors including:

- restructuring; • business portfolio changes; • traditional infrastructure-business reduction; • acquisitions and divestments.[10]

Evidence-safe conclusion:

Deutsche Telekom is undergoing broader workforce restructuring and digital transformation, but the reviewed evidence does not isolate RAN Guardian as a cause of headcount change.

This distinction is essential.

Otherwise a company-wide trend could be mistakenly converted into an AI-agent employment effect.

INDUSTRY CONTEXT

External Benchmark: Autonomous Networks Are Becoming an Industry Operating Model

Deutsche Telekom is not alone in moving from automation toward network autonomy.

Industry standards and frameworks increasingly describe autonomous networks as a progression from manual operation to closed-loop AI-driven operation.

01

TM Forum: levels of autonomy

TM Forum's Autonomous Networks framework uses a six-level taxonomy from fully manual operation toward full autonomy.[12]

Its description of higher levels includes capabilities such as:

- intent-driven management; • predictive analysis; • closed-loop control; • continuous learning; • service and customer-experience optimization.[12]

Level 4, for example, introduces intent-driven predictive decision-making and closed-loop management based on AI modelling and continuous learning.[12]

02

Boundary

AIGW should not label RAN Guardian “Level 4” unless Deutsche Telekom or an authoritative assessment does so.

The value of the TM Forum framework is comparative:

RAN Guardian exhibits several characteristics associated with higher-autonomy operating models, including closed-loop execution and predictive analysis.

That is different from a formal maturity certification.

03

ETSI: autonomy reduces manual workflows but does not remove human strategic responsibility

ETSI's work on autonomous networks notes that higher autonomy levels involve fewer manual workflows.[13]

Its 2026 discussion of AI agents in core networks describes agents as entities capable of:

- context acquisition; • reasoning; • decision-making; • execution; • collaboration.[14]

ETSI also warns that autonomous networks raise challenges around:

- reliability; • transparency; • interoperability; • security; • human oversight.[15]

Its assessment is especially relevant to this case:

04

ITU: autonomous operations need a control framework

ITU-T Recommendation M.3080, approved in March 2026, provides a framework for AI-enhanced telecom operation and management.[16]

It highlights requirements including:

- efficiency; • quality assurance; • cost management; • security assurance; • AI pipelines; • model management; • data management; • policy.

It also points to security controls such as:

- authentication; • access control; • non-repudiation; • audit trail; • alarm reporting.[16]

This matters because autonomy changes not only the workflow but the control plane around the workflow.

OPERATING MODEL

The New Scarce Resource: Operational Context

Autonomous systems need more than telemetry.

They need context.

TM Forum's 2026 Context Management work is useful here because it identifies a core transition inside network operations.[7]

Historically, context often existed inside human expertise.

An experienced network engineer might know:

- which alarm is noise; • what happened during a similar event last year; • what downstream system depends on a given site; • how a customer-impact pattern changes the urgency; • which configuration is safe under unusual conditions.

That knowledge may be partly tacit.

AI-native operations create pressure to make it explicit.

01

Context becomes infrastructure

The future operating stack increasingly needs structured access to:

- network state; • history; • incident data; • customer impact; • policies; • constraints; • topology; • operational intent.

The RAN Guardian's first agent already demonstrates this logic in a limited way by turning public-event information into machine-usable operational context. [1]

MINDR expands the challenge by correlating network signals across domains.[2]

02

AIGW interpretation: Context Transfer

This creates a workforce transition analogous to what LONGi and JD revealed in different sectors.

At LONGi:

inspector knowledge becomes AI-training and verification input.

At JD:

customer-service knowledge becomes agent-configuration input.

In network operations:

operational context that once lived partly in expert judgment increasingly has to become machine-readable, governed system context.

The question is not only whether AI can perform the action.

It is whether the organization can correctly encode the context required for safe action.

AIGW ANALYTICAL FRAMEWORK

AIGW Agentic Operations Ladder

FIGURE 3.1

AIGW framework

Six analytical levels of delegated operation

Level	Operational capability
Level A — Observe	AI monitors telemetry and gathers information. Human interprets and acts. [1] [2] [12]
Level B — Diagnose	AI identifies anomalies, correlates evidence or suggests root causes. Human decides what to do. [1] [2] [12]
Level C — Recommend	AI proposes an action. Human approves and executes. [1] [2] [12]
Level D — Execute within bounds	AI performs a pre-authorized action under defined conditions. Human supervises or reviews. [1] [2] [12]
Level E — Closed-loop adapt	AI monitors, diagnoses, acts, validates outcome and adjusts without requiring human execution of each step. [1] [2] [12]
Level F — Cross-domain orchestration	Multiple agents coordinate across systems and domains to optimize an end-to-end service or objective. [1] [2] [12]

Source: AIGW framework, contextualized by company-reported use-case evidence and the TM Forum autonomy taxonomy [\[1\]](#) [\[2\]](#) [\[12\]](#)

Note: These are analytical levels; they do not certify Deutsche Telekom at a formal TM Forum level.

AIGW ANALYTICAL FRAMEWORK

AIGW Agentic Control Plane

FIGURE 3.2

AIGW framework

Identity, access, boundaries and approval

Control dimension	Question the operating model must answer
1. Identity	Every agent should have a clear identity. Question: Which system or agent performed this action? [3] [4] [16] [17] [18]
2. Access rights	Agents should have defined permissions. Question: Which systems, data and actions can this agent access? [3] [4] [16] [17] [18]
3. Action boundaries	Agents need scope limits. Question: What is the maximum action this agent can take autonomously? [3] [4] [16] [17] [18]
4. Risk classification	Actions should be differentiated by consequence. Question: Is this a reversible low-risk optimization or a critical change? [3] [4] [16] [17] [18]
5. Approval rules	Some actions may proceed autonomously; others may require human approval. Question: At what risk threshold does autonomy stop? [3] [4] [16] [17] [18]

Source: AIGW control framework, informed by Deutsche Telekom and standards sources [\[3\]](#) [\[4\]](#) [\[16\]](#) [\[17\]](#) [\[18\]](#) [\[3\]](#) [\[4\]](#) [\[16\]](#) [\[17\]](#) [\[18\]](#)

Note: The enterprise governance material does not establish RAN Guardian-specific controls.

AIGW ANALYTICAL FRAMEWORK

Agentic Control Plane: Oversight and Accountability

FIGURE 3.3

AIGW framework

Logging, limits, monitoring and accountability

Control dimension	Question the operating model must answer
6. Logging	Actions should be recorded. Question: Can the organization reconstruct what happened and why? [3] [4] [16] [17] [18]
7. Limits	Cost, usage or action limits can constrain agents. Question: What prevents runaway or excessive execution? [3] [4] [16] [17] [18]
8. Monitoring and intervention	Humans need visibility and a way to intervene. Question: Who watches the agent landscape, and how quickly can they stop it? [3] [4] [16] [17] [18]
9. Accountability	Responsibility remains organizational. Question: Which person or function owns the outcome? [3] [4] [16] [17] [18]

Source: AIGW control framework, informed by Deutsche Telekom and standards sources [3] [4] [16] [17] [18] [3] [4] [16] [17] [18]

Note: The enterprise governance material does not establish RAN Guardian-specific controls.

FRAMEWORK APPLICATION

Applying autonomy levels and controls

01

Where RAN Guardian and MINDR sit

The Deutsche Telekom case suggests a useful ladder for distinguishing levels of AI participation in operational work.

Where RAN Guardian sits:

Public evidence supports:

- Observe; • Diagnose; • Recommend; • Execute; • document outcome.

The exact degree of autonomous post-action validation is less fully described.

Where MINDR is headed:

MINDR is designed toward:

- cross-domain diagnosis; • coordinated multi-agent operation; • controlled remediation; • end-to-end service orientation.[2]

This is closer to the upper end of the ladder as a design direction.

Why the ladder matters:

The workforce and governance implications increase as the system moves upward.

At Level A, the human still operates the process.

At Level D or above, the human increasingly operates the rules governing the process.

02

Why action-level controls matter

Once an AI system can act, organizations need a governance layer around action.

The September 2026 Deutsche Telekom agent-governance material provides a useful set of building blocks.[3][4]

AIGW organizes them into an Agentic Control Plane.

Why this is more useful than “human in the loop”:

Human involvement is not binary.

A mature control system can permit autonomous low-risk action while requiring human approval for higher-risk changes.

The central design variable is therefore: risk-adjusted autonomy.

03

Human oversight and strategic direction

routine operations can become increasingly autonomous while the human role remains important for oversight and strategic direction.[15]

04

How network expertise may shift

AIGW implication:

The network expert may move from:

operator of individual network actions toward:

designer, supervisor and assurer of autonomous network behavior.

Again, that is a role implication, not a documented Deutsche Telekom workforce outcome.

OPERATIONAL ECONOMICS

Verification After Action: The Missing Half of Autonomous Productivity

Automation is often evaluated by time to action.

Autonomous operations need another metric:

time and effort required to verify that the action was safe and effective.

A system can act in one minute.

If engineers then spend an hour investigating whether the action caused unintended consequences, the productivity story changes.

The public RAN Guardian material says actions and outcomes are documented, and later MINDR material emphasizes explainable remediation.[1][2]

But detailed evidence on the human verification burden is not published.

01

AIGW concept: Autonomous Operations Economics

A more complete value equation is:

time saved in detection/diagnosis/execution – human monitoring + verification + exception burden – governance / remediation cost = reliable operational gain

This extends the Verification Economics concept from high-stakes professional work into infrastructure operations.

02

What the measures establish

The strongest RAN Guardian evidence is at:

- deployment; • workflow; • autonomous-action levels.

The weakest evidence is at:

- role; • worker; • staffing levels.

This is precisely why the case belongs in a workforce transition tracker without inventing a workforce response.

WORKFORCE IMPLICATION

Workforce Implication: From Network Operator to Autonomy Supervisor?

The most tempting interpretation of the case would be:

AI agents will replace network operations staff.

The evidence does not support that conclusion.

A more defensible interpretation is that the task composition of network operations is changing.

Tasks likely to shrink in relative human share:

- repeated event discovery; • routine telemetry correlation; • manual first-pass diagnosis; • standard corrective actions; • repetitive documentation.

Network-domain expertise:

An agent needs meaningful policies and constraints.

People need to understand what safe behavior looks like.

AI orchestration:

Network teams increasingly need to understand:

- what agents are doing; • how agents coordinate; • when they fail; • how to constrain them.

Risk judgment:

A standard configuration change and a high-impact cross-domain action should not be governed identically.

Assurance:

Operators need to verify whether autonomous action preserved:

- service quality; • resilience; • security; • customer outcomes.

Incident exception handling:

Novel failures become more important as routine incidents are automated.

Policy and intent design:

If the AI executes the action, humans increasingly define:

- objectives; • constraints; • priorities; • risk tolerances.

MEASURED EVIDENCE

Measured evidence: operational outcomes

FIGURE 4.1

reported evidence

Reported operating outcomes

Evidence	Reported result	Analytical level and boundary
RAN Guardian deployment	Live in Germany from Nov 2025	Deployment — First-party ^[1]
Selected process duration	~1 hour → few minutes	Workflow — Representative process claim ^[1]
Major-event management	Hours → ~1 minute	Workflow — Company-reported >95% improvement ^[2]
Autonomous remediations	>100 in first month	Agent action — Christmas-market context ^[2]
2026 events identified	~237,000	Context / scale — Definition differs from narrower event counts ^[2]
Carnival events	~130	Use-case scale — Each >10,000 expected attendees ^[2]
Mobile sites pre-checked	611	Network operations — Carnival example ^[2]
Sites adjusted at peak	5	Autonomous optimization — Carnival example ^[2]

Source: Company-reported operating evidence ^[1] ^[2] ^[1] ^[2]

Note: The event counts have different scopes.

MEASURED EVIDENCE

Measured evidence: workforce context

FIGURE 4.2

reported evidence

Enterprise-wide workforce and learning indicators

Evidence	Reported result	Analytical level and boundary
Enterprise AI training	30,000 employees in 2025	Workforce context — Not RAN Guardian-specific ^[6]
AI-related LEX participation	52,000 participations	Workforce context — Not unique employees; not use-case-specific ^[6]
Skill-management participants	~70,000	Workforce context — Enterprise-wide ^[6]
Growth Hub users at launch	110,000	Talent infrastructure — Enterprise-wide ^[6]

Source: Deutsche Telekom workforce reporting [6] [6]

Note: These enterprise-wide figures do not establish a RAN Guardian-specific workforce response.

CROSS-CASE COMPARISON

Human authority across two high-stakes contexts

Kaiser Permanente and Deutsche Telekom make an unusually useful pair.

Kaiser Permanente ambient documentation:

AI:

- listens; • transcribes; • drafts.

Human physician:

- reviews; • edits; • finalizes documentation; • retains clinical decision authority.

The AI is deliberately bounded away from the highest-stakes decision.

Deutsche Telekom RAN Guardian:

AI:

- monitors; • analyzes; • recommends; • executes corrective network changes.

Human network expert:

- interacts with the system; • but detailed approval/override boundaries are not publicly specified in the use-case source.

The lesson:

“Human oversight” cannot be evaluated by asking whether a person is somewhere in the process.

The relevant comparison is:

Where does autonomous authority stop?

Kaiser draws the stop line before the clinical decision.

RAN Guardian publicly demonstrates autonomy extending into operational execution.

That does not make one model “better.”

They involve different risk structures.

It does show why governance must be use-case-specific.

DECISION IMPLICATIONS

Decision Implications

01

For CTOs and network leaders

Map autonomy by action, not by system.

One agent may safely execute a reversible capacity adjustment while a different action requires human confirmation.

Do not give an entire system one binary label: autonomous / not autonomous.

02

For CIOs and AI governance leaders

Build the control plane before scale makes it unavoidable.

At small scale, humans can manually monitor a few agents.

At large scale, governance needs:

- identity; • permissions; • logs; • risk tiers; • approval rules; • limits; • rollback; • accountability.

03

For CHROs

Do not wait for headcount change to begin workforce design.

If operational execution is becoming autonomous, ask now:

- which tasks shrink; • which expert knowledge must be retained; • which new supervisory capabilities are required; • how network careers change when routine operation is delegated.

04

For L&D leaders

Train for autonomy supervision, not generic AI fluency alone.

Relevant capability may include:

- network-domain expertise; • agent behavior; • model limits; • policy setting; • risk classification; • exception handling; • assurance; • incident recovery.

05

For network engineers

The durable advantage is not manually executing every known remediation.

It is understanding:

- when an autonomous action is appropriate; • when it is unsafe; • what context the system is missing; • how to recover when automated behavior fails.

06

For regulators and standards bodies

Autonomous infrastructure needs auditability.

Useful requirements include:

- agent identity; • access controls; • action logs; • traceability; • risk-based approval; • rollback; • accountability.

EVIDENCE BOUNDARIES

What the evidence establishes — and does not establish

01

What the Evidence Establishes

The reviewed evidence establishes that:

- Deutsche Telekom developed RAN Guardian with Google Cloud and tested/verified it before live use.
- RAN Guardian went live in German mobile-network management in November 2025.
- It operates as a multi-agent system.
- It identifies events, analyzes network conditions, detects anomalies and recommends optimization actions.
- It can autonomously implement corrective actions including network-resource reallocation and configuration adjustment.
- It documents actions and outcomes.
- Deutsche Telekom initially reported selected processes falling from around an hour to a few minutes.
- Follow-on reporting said major-event management fell from hours to around one minute, more than 95%.
- More than 100 autonomous remediation actions were triggered during the first month around Christmas-market events.
- RAN Guardian identified roughly 237,000 events for 2026 in February reporting.
- It pre-checked hundreds of mobile sites in the Carnival example and adjusted sites under peak load.
- Deutsche Telekom began scaling RAN Guardian beyond Germany, starting with the Czech Republic and Croatia.
- Deutsche Telekom announced MINDR as a broader multi-domain autonomous-diagnostics and remediation architecture.
- Later enterprise agent-governance material from Deutsche Telekom explicitly discusses agent identity, permissions, logging, risk-based human intervention and human approval for particularly critical actions.
- Deutsche Telekom has a large enterprise AI-skilling and skills-management infrastructure.
- Industry bodies including TM Forum, ETSI and ITU are formalizing autonomous-network and AI-operations frameworks.

02

What the Evidence Does Not Establish

The reviewed evidence does not establish:

- a specific network-operations headcount reduction caused by RAN Guardian;
- layoffs caused by the system;
- reduced hiring caused by the system;
- redeployment of network experts;
- a RAN Guardian-specific reskilling programme;
- a formally observed capability shift among network staff;
- the percentage of all network-management work automated;
- the percentage of RAN Guardian actions requiring human approval;
- whether every configuration change is reversible automatically;
- whether engineers can veto every proposed action;
- whether every autonomous action can be overridden in real time;
- the complete rollback architecture;
- the confidence thresholds used for autonomous action;
- the exact division of accountability between engineering teams, product owners and platform governance;
- that RAN Guardian uses Deutsche Telekom's later T-AI Agentic Hub governance controls;
- that RAN Guardian has been independently certified at a specific TM Forum Autonomous Network level;
- that MINDR had the same production maturity as RAN Guardian at the time of the reviewed announcements;
- that company-wide headcount restructuring was caused by autonomous network operations;
- that faster event management directly reduced operating cost by a specified amount;
- that autonomous remediation improved customer outcomes by a quantified amount independent of other network investments.

RESEARCH AGENDA

Research Gaps to Watch

01

Research Gaps to Watch

1. Which RAN Guardian actions are fully autonomous?; 2. Which require human pre-approval?; 3. What risk taxonomy is used for actions?; 4. Can network experts veto an action before execution?; 5. Can they override or roll back an action after execution?; 6. Is rollback automated?; 7. What confidence threshold determines autonomous action?; 8. What happens when agents disagree?; 9. How are novel/unseen network conditions handled?; 10. What action logs are retained?; 11. Which organizational role is accountable for autonomous network changes?; 12. How often do humans intervene?; 13. What is the false-positive rate for event/network-risk identification?; 14. What is the rate of unnecessary remediation?; 15. What is the rate of failed remediation?; 16. How much human verification is required after an action?; 17. How have network-expert task portfolios changed since launch?; 18. Has staffing changed in network operations?; 19. Has on-call work changed?; 20. Has incident-response workload changed?; 21. What new skills are formally required?; 22. Is there a network-specific AI-agent training curriculum?; 23. Are employees certified before supervising autonomous actions?; 24. How are works councils/social partners involved in role redesign?; 25. What parts of tacit NOC knowledge are being encoded into machine-readable context?; 26. Does the system change career entry for junior network engineers?; 27. How portable are AI-orchestration skills across telecom employers?; 28. What is the production status of MINDR across RAN, transport and core by end-2026?; 29. How is cross-domain blast radius controlled?; 30. Can independent evaluation confirm the >95% event-management improvement?; 31. What customer-experience improvement can be attributed specifically to agentic remediation?; 32. How will Deutsche Telekom map its systems to emerging TM Forum/ETSI/ITU autonomy frameworks?

METHOD AND SOURCE QUALITY

Method and Source Quality

This Deep Dive separates six evidence layers.

01

A. Production use-case evidence

Deutsche Telekom's November 2025 RAN Guardian release is the canonical source for:

- production deployment; • multi-agent workflow; • autonomous action; • one-hour-to-minutes process compression.[1]

It is first-party evidence.

It is strong for establishing what the company deployed and what the system is authorized to do.

It is weaker as independent validation of performance.

02

B. Follow-on operating evidence

Deutsche Telekom's February 2026 MINDR/RAN Guardian reporting adds:

- >100 autonomous remediations; • >95% event-management time improvement; • 237,000 identified events; • specific Carnival site activity; • cross-country scale-out.[2]

This materially strengthens the operational case.

It remains company-reported.

03

C. Follow-on architecture evidence

MINDR sources establish the direction toward:

- end-to-end service operations; • RAN + transport + core; • multi-agent coordination; • predictive remediation.[2][5]

They should not be treated as equivalent to mature production outcome evidence unless later sources confirm deployment.

04

D. Governance evidence

Deutsche Telekom's September 2026 agent-governance materials provide unusually useful institutional statements about:

- identity; • access rights; • logging; • boundaries; • risk-based human intervention; • human approval for critical actions; • retained organizational responsibility.[3][4]

They are broader enterprise-agent governance evidence, not RAN Guardian-specific controls.

05

E. Workforce evidence

The CR Report and HR Factbook establish:

- large-scale AI learning; • skills management; • internal talent infrastructure; • broader workforce restructuring.[6][10][11]

They do not establish the workforce response to RAN Guardian.

06

F. External institutional benchmarks

TM Forum, ETSI and ITU sources are used to frame:

- autonomy maturity; • context management; • closed-loop operation; • security; • human oversight; • emerging industry control requirements.[7][12][14][15][16]

They are not Deutsche Telekom outcome evidence.

AIGW INTERPRETATION

AIGW Interpretation: The Unit of Governance Is the Action

The deepest lesson from the case is methodological.

Organizations often talk about whether an AI system is autonomous.

That is too coarse.

The better unit of governance is the action.

A single system may:

- monitor autonomously; • diagnose autonomously; • recommend autonomously; • execute one class of action autonomously; • require approval for another class; • prohibit a third class entirely.

This leads to a more useful model:

Action × Risk × Reversibility × Human Authority

For each possible action, ask:

- What is the consequence if wrong? • Is the action reversible? • How quickly can harm propagate? • How confident is the system? • Is human approval required? • Who can stop or reverse it? • Who is accountable?

01

AIGW operating rule

Do not assign autonomy to a system. Assign autonomy to specific actions under specific conditions.

That principle is likely to become increasingly important as agentic AI moves into:

- networks; • finance; • industrial systems; • public services; • healthcare operations.

02

Workforce evidence boundary

RAN Guardian makes AI orchestration, assurance and exception management increasingly plausible future capabilities for network experts, but Deutsche Telekom has not publicly documented a RAN Guardian-specific role redesign or workforce transition at the level required to classify those shifts as observed outcomes.

CONCLUSION

Conclusion

Deutsche Telekom's RAN Guardian Agent marks a meaningful shift in enterprise AI.

The system is not only assisting network experts.

It is performing part of the operational loop itself.

It gathers context, analyzes network conditions, identifies anomalies, recommends actions and can autonomously change network resources or configuration.

The process improvement is material.

Tasks that previously took around an hour were reported to take minutes. Follow-on reporting described major-event management falling to around one minute and more than 100 autonomous remediations in the system's first month of operation.

But the most important feature of the case is not speed.

It is delegated authority.

Once AI is allowed to act on infrastructure, governance must become more precise.

The relevant questions are no longer only:

- Is the model accurate? • Does it save time?

They become:

- What may the agent change? • Under what conditions? • Which actions require approval? • Who can override it? • How is the action verified? • Who remains accountable?

The RAN Guardian public record answers the first question – what the agent can do – much more clearly than it answers the full set of human-control questions.

That is why the case is valuable.

CONCLUSION

It reveals the Authority Gap that can emerge as production AI becomes agentic.

Deutsche Telekom's later enterprise agent-governance materials show the direction of travel: clear identities, access rights, logged actions, risk-based intervention and retained human/organizational responsibility.

The next challenge is to connect those governance principles to use-case-level evidence with the same specificity that companies now use to report productivity.

From a workforce perspective, the case should not be misread as a network-engineer replacement story.

The evidence supports workflow redesign and task compression.

The workforce outcome remains open.

What is changing already is the nature of expertise.

As routine detection, diagnosis and remediation become autonomous, the scarce human contribution is likely to move toward:

- defining intent; • governing permissions; • handling exceptions; • verifying outcomes; • supervising agents; • designing resilient systems.

The network expert does not disappear from the operating model.

The locus of human value moves upward — from executing every action toward governing the conditions under which actions may execute themselves.

That is the deeper transition.

And it gives us a more precise question for the agentic era:

When AI can act, what must humans still control?

REFERENCES

Notes and sources

1. Deutsche Telekom. *AI agents for mobile network*. 11 November 2025. Company / first-party production use-case evidence.
Supports: live RAN Guardian deployment; multi-agent workflow; event discovery; network-capacity analysis; autonomous corrective action; resource reallocation/configuration adjustment; outcome logging; approximately one-hour-to-minutes process compression; interaction between network experts and AI.
Boundary: does not specify full human approval/veto/override architecture or workforce consequences.
<https://www.telekom.com/en/newsroom/latest-updates/media-information/2025/11/deutsche-telekom-ai-agents-for-mobile-network>
2. Deutsche Telekom. *Deutsche Telekom and Google Cloud Collaborate for Superior Network Experience with Agentic AI*. 25 February 2026. Company / first-party follow-on operational evidence.
Supports: >100 autonomously triggered remediation actions in first month; major-event management from hours to around one minute / >95%; ~237,000 events identified for 2026; Carnival example; 611 sites pre-checked; five sites optimized; Czech Republic/Croatia scaling; MINDR architecture and planned production releases.
Boundary: MINDR announcement/implementation direction should not be treated as having identical production maturity to RAN Guardian.
<https://www.telekom.com/en/newsroom/latest-updates/media-information/2026/2/mindr-ai-agents-in-telekom-network>
3. Deutsche Telekom. *How much should an AI agent be allowed to decide on its own?*. 8 September 2026. Company / expert governance interview.
Supports: responsibility cannot be delegated to an agent; use-case-specific autonomy; human testing before release; risk-based intervention; critical actions can require human approval; cost and usage limits.
Boundary: enterprise agent-governance model; not direct proof of RAN Guardian's internal controls.
<https://www.telekom.com/en/newsroom/latest-updates/news/2026/09/how-much-should-an-ai-agent-be-allowed-to-decide-on-its-own>
4. Deutsche Telekom. *Many AI agents, one command center: Telekom gives you a clear overview*. 8 September 2026. Company / enterprise-agent governance product reporting.
Supports: agent identity, responsibilities, access-control and governance framing; T-AI Agent Platform and T-AI Agentic Hub.
Boundary: commercial/enterprise platform evidence, not RAN Guardian-specific implementation evidence.
<https://www.telekom.com/en/newsroom/latest-updates/media-information/2026/09/many-ai-agents-one-command-center-telekom-gives-you-a-clear-ov>
5. Deutsche Telekom. *MWC 2026: World premiere of AI-powered call assistant / MINDR section*. March 2026. Company / technology roadmap evidence.
Supports: MINDR as evolution of RAN Guardian; cross-domain RAN/transport/core operation; predictive anomaly detection; root-cause analysis; autonomous remediation; initial deployments planned in 2026.
<https://www.telekom.com/en/newsroom/latest-updates/media-information/2026/3/mwc-2026-world-premiere-of-ai-powered-call-assistant>
6. Deutsche Telekom. *CR Report 2025: Employee development*. n.d.. Company sustainability / workforce report.
Supports: 30,000 employees trained in AI-related topics in 2025; 52,000 AI-related LEX participations; AI-use targets; digital expert share; ~70,000 completing skill management; 110,000 Growth Hub users; transfer-management and learning infrastructure.
Boundary: enterprise workforce context; not RAN Guardian-specific reskilling evidence.
<https://report.telekom.com/cr-report/2025/social/employee-development.html>

REFERENCES

7. TM Forum. *Context Management for AI-Native Operations*. 2026. Industry standards / operational benchmark.
Supports: operational context currently created/interpreted by NOC engineers; need to make context explicit, governed and available to AI agents and humans in AI-native operations.
Boundary: external benchmark only.
<https://www.tmforum.org/resources/guidebook/ig1547-context-management-for-ai-native-operations-v1-1-0/>
8. Deutsche Telekom. *The Telekom network 2026: toward autonomous networks*. January 2026. Company / network operations update.
Supports: RAN Guardian operating since November; autonomous corrections within minutes; around 240,000 events identified/classified for 2026; roughly 100x previous event-identification/adjustment scale in the relevant process.
Boundary: event definition may differ from other later counts.
<https://www.telekom.com/de/newsroom/aktuelles/medieninformationen/2026/01/netztag-2026>
9. Deutsche Telekom. *Festival Summer 2026: AI-supported network optimization*. 17 September 2026. Company / operational follow-on reporting.
Supports: RAN Guardian used to detect public events, assess expected load and derive optimization; narrower 1,750 open-air-event context.
Boundary: should not be treated as directly comparable with the ~237k / ~240k broader event counts.
<https://www.telekom.com/de/newsroom/aktuelles/medieninformationen/2026/09/datenverkehr-im-festival-sommer-2026>
10. Deutsche Telekom. *HR Factbook 2025: Employees by operating segment*. n.d.. Company workforce report.
Supports: Group and segment headcount; Germany operating segment 57,303 FTEs in 2024 to 55,089 in 2025; broader restructuring context.
Boundary: no causal link to RAN Guardian.
<https://report.telekom.com/hr-factbook/2025/our-employees/employees-by-operating-segment.html>
11. Deutsche Telekom. *Annual Report 2025 / Group employment data*. n.d.. Company annual report.
Supports: Germany employee totals and Group productivity/headcount context; socially responsible restructuring references.
Boundary: enterprise-level context only.
<https://report.telekom.com/annual-report-2025/management-report/development-of-business-in-the-group/results-of-operations-of-the-group.html>
12. TM Forum. *Autonomous Networks Mission / Autonomous Network levels taxonomy*. n.d.. Industry framework.
Supports: six-level autonomy taxonomy; higher-level characteristics including intent-driven predictive analysis, closed-loop management, AI modelling and continuous learning.
Boundary: does not establish Deutsche Telekom's formal AN level.
<https://www.tmforum.org/missions/autonomous-networks>
13. ETSI. *Driving optical networks towards higher Autonomous Network levels*. 24 June 2025. Standards-body technical context.
Supports: higher autonomous-network levels correspond to fewer manual workflows; operational tasks include intent translation, awareness, analysis, decision and execution.
<https://www.etsi.org/newsroom/blogs/entry/driving-the-optical-networks-towards-higher-autonomous-network-levels>

REFERENCES

14. ETSI. *AI agents: the new PAs of the core network?*. 15 May 2026. Standards-body technical context.
Supports: AI agents can perceive environment, acquire context, reason, decide, execute and collaborate; autonomous-agent role in future telecom networks.
<https://www.etsi.org/enjoy-magazine/articles/ai-agents-the-new-pas-of-the-core-network/>
15. ETSI. *AI in the evolution of Autonomous Networks*. n.d.. Standards-body white paper.
Supports: reliability, transparency and interoperability challenges; human element remains important for oversight and strategic direction as routine operations become more autonomous.
https://www.etsi.org/images/files/ETSIWhitePapers/ETSI-WP-69-AI-intheevolutionofAutonomous_Networks.pdf
16. ITU-T Recommendation M.3080. *Framework of artificial intelligence enhanced telecommunication operation and management*. 29 March 2026. International technical standard.
Supports: AI-enhanced telecom operations framework; quality/cost/security assurance; model/data/policy functions; security controls including authentication, access control, non-repudiation and audit trail.
<https://www.itu.int/epublications/zh/publication/itu-t-m-3080-2026-03-framework-of-artificial-intelligence-enhanced-telecommunication-operation-and-management-aitom>
17. ITU-T. *Human-AI collaboration work item for intelligent network status awareness*. 2026. Emerging international standards work.
Supports: need for dynamic authority switching between AI agents and human experts based on confidence, unknown anomalies and risk; human-AI task allocation as an explicit mechanism.
Boundary: emerging benchmark, not Deutsche Telekom implementation evidence.
https://www.itu.int/ITU-T/workprog/wp_item.aspx?isn=23870
18. Deutsche Telekom. *Guidelines for Artificial Intelligence*. n.d.. Company AI governance principles.
Supports: binding internal AI guidelines, human-centric approach and explicit organizational responsibility.
Boundary: high-level governance principles; not use-case-specific control evidence.
<https://www.telekom.com/en/newsroom/topic-hubs/technology-and-innovation/ai/guidelines-for-artificial-intelligence>
19. Google Cloud. *Google Cloud and the rise of the agentic telco*. 2026. Technology-partner / industry context.
Supports: MINDR as a multi-agent cross-domain network-operations system; broader agentic-telco trajectory.
Boundary: partner source; not independent outcome evaluation.
<https://cloud.google.com/transform/google-cloud-rise-of-the-agentic-telco-MWC>