Responsible Agentic AI for Equitable and Sustainable Cloud Native Telecom

Bhavika Jalli

Conf42 Cloud Native 2026

The Stakes

Networks Are Becoming Autonomous Is Responsibility Keeping Pace?

5G → 6G Evolution

Containerised network functions and and distributed edge workloads now now manage spectrum, faults, traffic, traffic, and power in real time.

AI-Driven Automation

Reinforcement learning, digital twins, twins, and agentic orchestration are replacing human-in-the-loop decisions at decisions at scale.

Responsibility as a Cloud Native Concern

As networks become software-defined and agent-driven, governance must be designed in — not bolted on.

AGENDA

What We'll Cover Today

01

The Connectivity Divide

Why 2.6 billion people remain under- or unconnected and how AI allocation patterns can widen that gap.

02

Bias at Cloud Scale

How historically skewed training data propagates inequity through through cloud native orchestration layers.

03

The Sustainability Equation

Balancing real-time power optimisation against the computational cost computational cost of large-scale AI inference.

04

Integrated Governance Framework

Embedding fairness, transparency, accountability, and carbon reporting into autonomous network operations.

Chapter 1

A World Still Divided by Connectivity

Access is not evenly distributed and for billions, the difference shapes how they how they learn, work, and participate online.

Connected: 4.5 b

Enjoy reliable, high-speed internet that supports daily digital life, from streaming and collaboration to essential services.

Under-connected: 1.5 b b

Navigate a fragile connection that slows slows downloads, disrupts calls, and makes everyday tasks tasks inconsistent.

Unconnected:1 b 1 b

Remain cut off from the online world, unable to access the tools, information, and opportunities it provides.

CHAPTER 2

Bias at Cloud Scale

When Automation Inherits Inequity

AI resource allocation models trained on historical investment data mirror decades of decades of uneven infrastructure spending. In cloud-native systems, those biases do biases do not remain confined they spread across distributed orchestration layers at layers at machine speed.

How Bias Scales in Cloud Native Orchestration

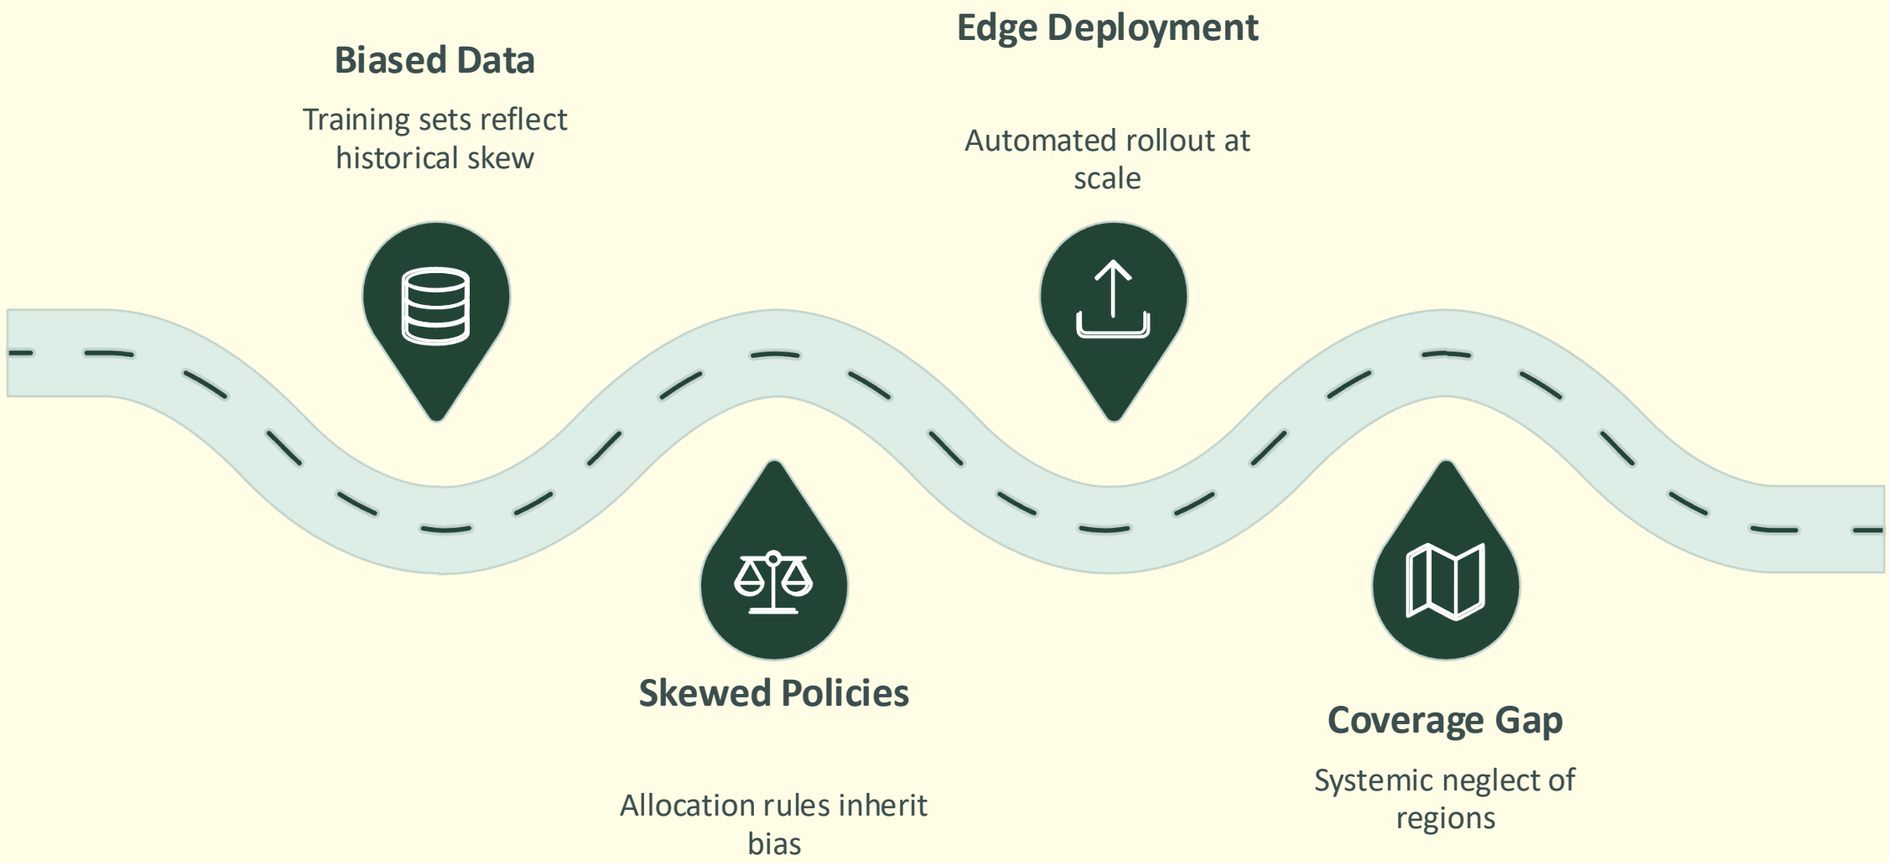

Each layer of automation amplifies upstream bias. A skewed reward signal in a reinforcement learning agent can translate at container orchestration speed into orchestration speed into systematic neglect of underserved regions across hundreds of distributed edge nodes.

Bias in Action: A Hypothetical Hypothetical Scenario

The Setup

A national operator trains an RL agent on 10 years of investment data to automate base station capacity planning across 500 edge nodes.

What Goes Wrong

Rural regions with historically low usage get systematically deprioritised. The agent interprets low demand as weak signal, not inadequate infrastructure. Coverage gaps widen at machine speed across 200+ nodes before any human notices.

The Fix

Coverage equity thresholds added as hard RL constraints. Geographic fairness audits at each deployment gate. Result: rural coverage improves 23% while urban efficiency gains are preserved.

CHAPTER 3

The Sustainability Equation

Reinforcement learning can cut idle base station energy consumption significantly. significantly. But large-scale model training and continuous edge inference carry inference carry their own substantial footprint. The net sustainability outcome outcome depends entirely on whether optimisation gains outpace the operational operational cost of AI-driven cloud workloads.

Two Forces in Tension

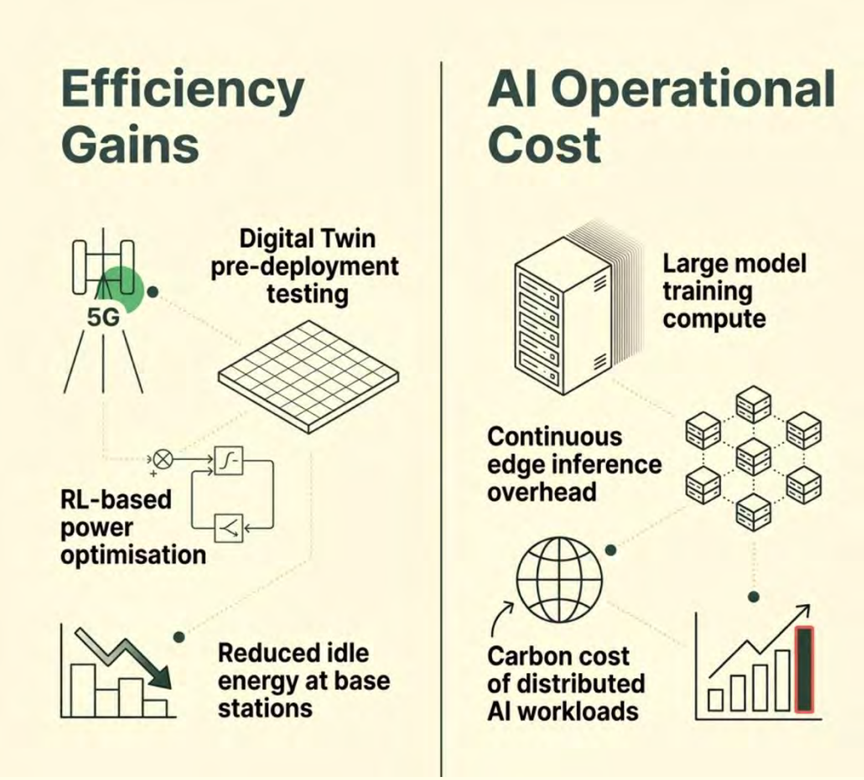

Optimise or Offset?

Digital twin simulations allow operators to validate efficiency strategies before live deployment avoiding costly trial-and-error at the network edge. The challenge is ensuring the AI infrastructure enabling these gains does not negate them through unchecked compute consumption.

CHAPTER 4

Governance by Design

Governance must be architected in from the ground up — not retrofitted after deployment. Fairness checks, carbon tracking, policy enforcement, and accountability hooks belong in the delivery stack itself.

An Integrated Governance Framework

Transparency & Audit

Explainable AI decisions logged at every every orchestration layer, with immutable immutable audit trails accessible to regulators and operators.

Fairness Evaluation

Allocation assessments that test models against demographic and geographic equity metrics before and after deployment.

Accountability Mapping

Clear responsibility chains across cloud cloud native microservices every automated decision traceable to a responsible owner.

Carbon Reporting

Measurable sustainability tracking integrated into CI/CD pipelines, quantifying the carbon cost of every model update and inference inference workload.

Embedding Governance into Cloud Native Pipelines

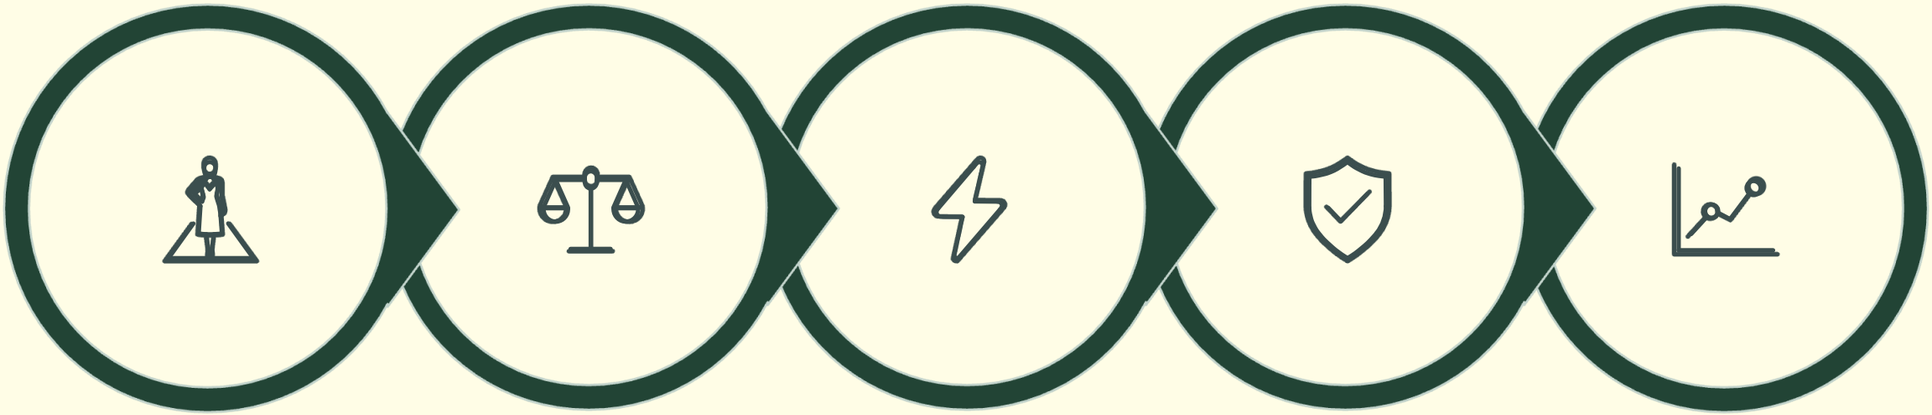

Governance is not a checkpoint at the end of a pipeline it is a continuous process woven into every stage, from training data curation through live network through live network operation and post-deployment observability.

Fairness as a Technical Specification

Designing Equity into Automation Workflows

Fairness must be operationalised as measurable constraints in the in the orchestration layer not aspirational policy language.

- Coverage equity thresholds as hard RL constraints
- Geographic distribution audits at deployment gates
- Demographic parity checks pre- and post-rollout

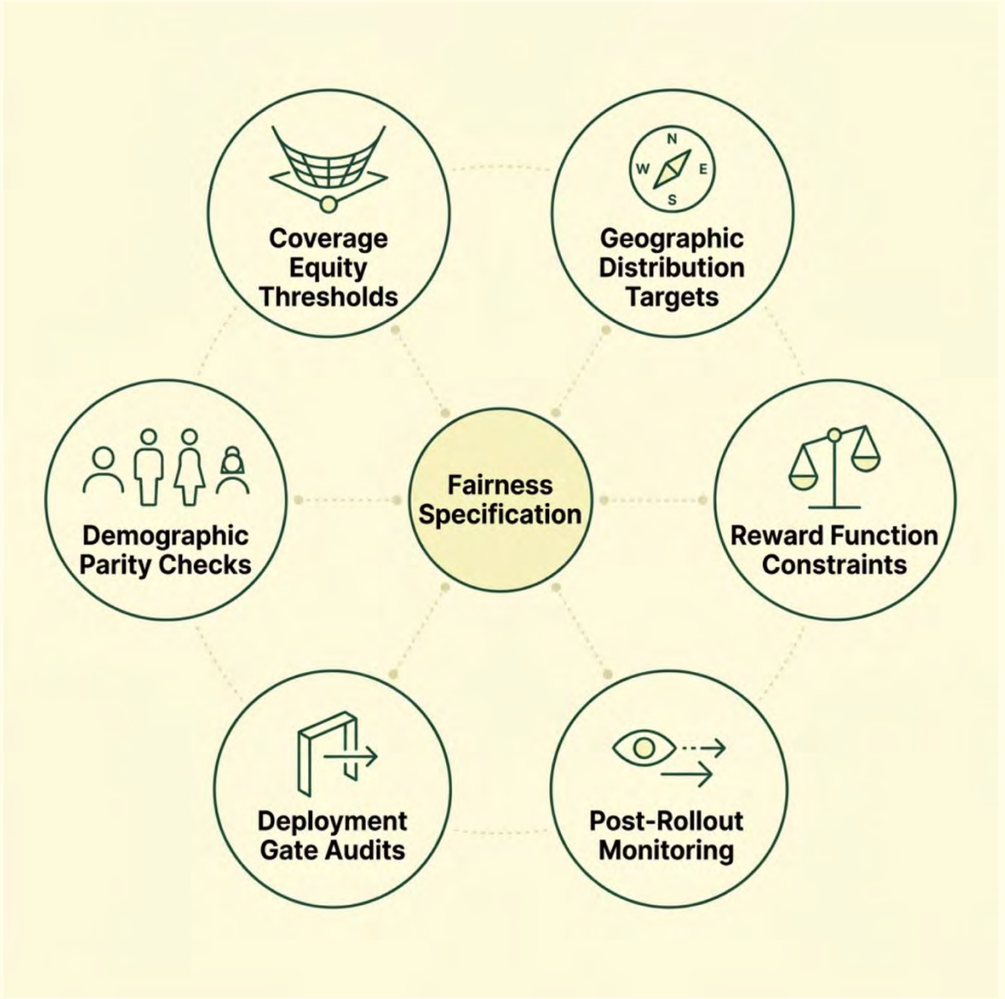

Responsible AI Enables Not Constrains Innovation

Scalable Societal Value

Governance can align fairness and carbon accountability with measurable measurable network gains.

Regulatory Readiness

Built-in auditability helps operators adapt to regulation faster and at lower at lower cost.

Trust at Network Scale

Transparent AI strengthens trust with regulators, communities, and enterprise customers.

Key Takeaways

Bias scales with automation

Historically skewed training data propagates inequity at machine speed speed across cloud native orchestration orchestration fairness must be a hard hard constraint, not a soft guideline.

Sustainability requires net accounting

Power optimisation gains are real, but but only meaningful if the AI footprint footprint enabling them is measured, measured, managed, and reported transparently.

Governance is an architecture architecture decision

Embed transparency, accountability, fairness evaluation, and carbon reporting reporting directly into cloud native pipelines before agents reach production. production.

Thank You!

Questions? Let’s connect.

Bhavika Jalli

LinkedIn: linkedin.com/in/bhavikajalli