

THE AUTONOMOUS GRID

Driving Zero-Touch Network Operations with Smart Networks

Amajit Gupta

Group CEO & Managing Director, Lightstorm

The scarce resource is not bandwidth, or compute. It is latency.

Capacity can be added. Compute can be scaled. A millisecond, once passed, is gone — and it takes the value of that transaction with it.



THE QUESTION HAS CHANGED

From “Is my network connected?”

To “Is my network responding at the speed
my application requires?”

■ 02 · THE ECONOMICS OF A MILLISECOND

The network has to decide before the workload loses value

10 ms

REAL-TIME VALUE

Interactive AI agents, fraud detection, industrial control, autonomous systems and computer vision all remain viable.

100+ ms

VALUE AT RISK

The same workload becomes slower, less responsive, or commercially unsuitable for the application it was built to serve.

THE AUTONOMOUS RESPONSE LOOP



Latency thresholds are illustrative and vary by application, architecture and geography.

■ 03 · THE EVIDENCE

AI is exposing the limits of manual network operations

What enterprises say is holding back AI-ready WANs — 269 IT professionals surveyed

42%

Workload distribution
Complexity across multiple sites

39%

WAN-edge latency
Between workloads and the WAN edge

36%

Traffic prioritization
Difficulty prioritizing AI traffic

33%

Network congestion
Network capacity is becoming a constraint

Sources: EMA via Network World; Ericsson Mobility Report

The issue is not that AI needs faster networks. It is that AI is exposing the structural limits of networks that are still operated by hand.

■ 04 · THE AUTOMATION GAP

The question is no longer whether networks can be automated

It is whether the industry can operationalise automation at scale.

18%

TODAY


of IT teams rate their network-automation strategy a complete success.

The intent is not in question.

Gartner expects enterprises to automate the majority of their network activity. Ambition is not the constraint — the operating model is.

The ambition is set. The operating model is the bottleneck — and that is precisely what zero-touch operations is designed to solve.

NaaS is not the future architecture — it is the transition



NaaS makes capacity programmable.
Autonomy makes performance programmable.


*Most enterprise WAN teams have already bought NaaS or are actively evaluating it, and providers now treat it as essential to enterprise services.
The open question is no longer whether to adopt NaaS — it is what comes after it.*

From manual reaction to a closed autonomous loop

TODAY — HUMAN-CENTRIC

- 1 A link degrades
- 2 Someone notices
- 3 Someone investigates
- 4 Someone decides what to do
- 5 Someone changes the configuration
- 6 Someone verifies the outcome

THE AUTONOMOUS GRID



Every action runs inside explicit human guardrails.

Four capabilities at the heart of the Autonomous Grid



Real-Time Telemetry

Continuous visibility into latency, congestion, capacity, route health and behaviour — not just up or down.



Intent-Based Orchestration

State the outcome you need. The network determines how to prioritise, protect latency, and stay resilient.



Programmable Path Control

The optimal path at 10:00 AM may not be optimal at 10:05. Paths are re-evaluated and moved as conditions change.



Resilient, Self-Correcting Design

The strongest network isn't one that never fails — it's one where failure never becomes a customer problem.

Lightstorm is building the foundations for smart network operations



Unified network
data model



Topology & service
visibility



Real-time
telemetry



Alarm
correlation



Predictive
fibre health



Evidence-based
RCA assistance



Programmable
path control



Controlled closed-loop
automation

Automation earns autonomy. Every closed loop runs inside explicit human guardrails — the network proposes, and the operator remains accountable.

Autonomy has a P&L

15–30%

Reduction in total network
OPEX

30–70%

Fewer troubleshooting tickets

55–80%

Lower network operations-
centre costs

30–40%

Faster mean time to repair

And the savings are not only operational



Energy has reached up to 5% of operator revenue —
an increasingly automatable cost line.

~\$20 m SAVED

in a single financial year by one large Indian mobile operator' s AI-
enabled energy-optimisation deployment across ~157,000 cell sites.

Automation assists engineers — it does not replace them

Those savings come from removing repetition, not from removing people. What changes is where engineering time goes.

MOVES TO SOFTWARE

The repetitive, high-frequency work

- Continuous detection and monitoring
- Alarm correlation and noise reduction
- First-pass triage of routine events
- Verification that a change worked

STAYS WITH ENGINEERS

The work that actually needs judgement

-  Network design and architecture
-  Resilience and failure strategy
-  Oversight and accountability
-  Judgement on novel exceptions

The scarce resource in this industry is engineering judgement. Automation should buy more time for it, not less.

■ 11 · THE REGIONAL OPPORTUNITY

From automated operations to a regional grid

Everything so far is what one operator can do alone. A true regional grid is a longer-term goal no single operator can build.

A LONGER-TERM, SHARED GOAL



Operators must stay in control of their own network

Network data, AI models, safety limits and manual override belong with the operator. A shared grid must not become dependence on any single external provider — autonomy should increase an operator's control, not transfer it.

Carriers

Subsea networks

Cloud providers

Data centres



Illustrative regional fabric — not a route map

Carrier-neutrality is what makes that grid usable: no customer forced into one cloud, one carrier or one route.

The differentiator is no longer connectivity.

It is the intelligence and autonomy of the connectivity.



**“The best network is the one
you never have to manage.”**

*Zero-touch does not mean zero technology, and it does not mean zero people.
It means zero unnecessary intervention.*