



Lyyli.ai

READY-TO-PUBLISH DRAFT

Power per bit: the metric that will govern network design in the AI supercycle

A publish-ready article in Nokia's voice.

Company	Nokia
Website	www.nokia.com
Length	about 730 words
Date	14 August 2026

AI traffic is growing faster than the energy budgets available to carry it. That makes energy per bit — not peak throughput alone — the constraint that decides which networks can scale. Here is how we approach efficiency across silicon, software and systems, and what operators should ask of their next build.

The shift: capacity is rising faster than power budgets

Every layer of the network is being asked to carry more. AI training and inference traffic moves between data centers, across metro and long-haul optical links, and increasingly out to the radio edge. Capacity demand is compounding.

Energy budgets are not. Power availability, grid connection timelines, cooling limits and operating costs now shape what can be deployed and when. For many networks, the binding constraint is no longer how many bits a platform can move, but how much power it consumes per bit moved.

That reframes the design question. Throughput alone is an incomplete specification. Throughput per watt, and the ability to hold that ratio as traffic scales, is the number that determines whether a network can absorb another decade of growth.

Efficiency is a silicon decision first

Efficiency cannot be added late. It is set largely by the components that process and transport traffic, which is why we invest in our own silicon across mobile, fixed and transport domains.

- **FP5** is our network processor for IP routing, designed to raise routed capacity while improving power per bit at scale.
- **PSE-6s** is our coherent optical engine, moving more traffic per wavelength over longer distances — reducing the amount of equipment needed to carry a given load.
- **Quillion** is our fiber access chipset family, supporting passive optical network (PON) generations on the same platform so operators avoid full hardware replacement at every step.
- In radio, we have reduced power consumption in 5G massive MIMO (mMIMO) by 50%.

The pattern is consistent: fewer components doing more work, with each generation improving the efficiency ratio rather than only the headline capacity figure.

Software decides how much of that hardware you actually use

Hardware sets the floor. Operations determine the real number. Networks are dimensioned for peak load, but they spend most of their hours well below it — and idle capacity still draws power.

This is where autonomous operations matter. Networks that sense, think and act can observe live conditions, decide where capacity is genuinely needed, and adjust in real time rather than by static configuration. Our **Autonomous Networks Suite** provides that closed-loop foundation, and **MantaRay SMO** brings AI-powered Service Management and Orchestration (SMO) to radio access network (RAN) operations, including energy savings management.

The operational value is straightforward. Energy features that require manual tuning are rarely applied at full scale, because the effort per site is too high. Features that run autonomously are applied everywhere, every hour. The difference between those two outcomes is usually larger than the difference between two hardware generations.

Evolution paths matter for the same reason. With **anyRAN**, operators can choose purpose-built, cloud or hybrid RAN deployments and combine them across a single network. That optionality allows efficiency upgrades to be phased against site economics instead of forcing a single, capital-intensive replacement cycle.

Our own numbers, stated with scope and baseline

We apply the same measurement discipline to our own operations.

- Nokia has committed to net-zero greenhouse gas (GHG) emissions across its value chain by 2040.
- We target a 50% reduction in absolute scope 1, 2 and 3 emissions by 2030 from a 2019 base year.
- Both targets are approved by the Science Based Targets initiative (SBTi).
- Renewable electricity accounted for 87% of our electricity use in 2024.

We separate two things deliberately. Our **footprint** is the emissions of our own operations and value chain. Our **handprint** is the emissions avoided by what our

technology enables in customer networks and industries. Footprint is what we reduce; handprint is what we help others reduce. Reporting them separately keeps both claims honest.

What to ask of your next network build

For telecommunication providers, AI and cloud providers, and mission-critical enterprises planning capacity for the AI era, four questions cut through most vendor material:

- 01** What is the power per bit at realistic load, not only at peak?
- 02** Which efficiency features run autonomously, and which require per-site engineering effort?
- 03** How many hardware replacements does the stated evolution path require between now and 2030?
- 04** What baseline year, scope and third-party validation sit behind the vendor's own emissions claims?

Efficiency is no longer a sustainability annex to the network plan. It is the capacity plan.

Explore our Autonomous Networks Suite. Read our approach to sustainability and science-based targets.

Why this topic?

It sits at the intersection of two themes Nokia leads with on its own site — the AI supercycle driving capacity demand, and quantified, SBTi-validated sustainability commitments — and connects them through named Nokia technologies (FP5, PSE-6s, Quillion, MantaRay SMO, anyRAN, Autonomous Networks Suite). It speaks to all three primary commercial audiences (telecommunication providers, AI and cloud providers, mission-critical enterprises) with an expert point of view rather than product promotion, matching the blog and insights register.

Suggested placement

Nokia Blog / Insights section (nokia.com/blog or nokia.com/insights), cross-linked from the Autonomous Networks hub, the AI and cloud providers page, and the Sustainability 'All about zero' page.

The content is yours

Use and edit this text freely in any channel. No strings attached, no attribution required.

NEXT STEP

How do you keep this consistent across the whole team?

A document helps, but real change happens when every writer hits the same voice every time. Let's spend 30 minutes on what that would look like for you.

[Book 30 minutes with Mikko](#)

Mikko Oksanen · CEO, Lyyli.ai · mikko@lyyli.ai



Lyyli.ai