

An Open Letter to My Fellow CEOs

Ok, let's be honest. How often do you think about your network architecture? Do you just delegate it to your CTO?

Who can blame you? Most of us have spent the last few years all-in on AI. We've debated models, talent, chips, and data. We've made commitments to our boards. We see AI as the defining technology of our time. Exciting stuff.

Networks, on the other hand, are boring. That is, until they become critical. Until they can't handle the job. Until they fail. And that leads to a hard truth: Many of us are trying to build the future on infrastructure that cannot support it.

For decades, networking has been in the background - like plumbing. Today, that framing is not only dated, but also detrimental. In an AI-driven enterprise, the network is more like the nervous system. It controls and coordinates. It determines how fast you can move, how much you spend, and whether your AI investments produce value.

This isn't a tweak. It's a hard reset. One that requires you to treat the network as an AI ROI lever: fund it, govern it, and measure it.

AI systems don't operate in a single place. They require constant data movement between clouds, data centers, and edge points. The new corporate workforce is comprised of AI agents and bots. They're proliferating rapidly, operating continuously, insatiably consuming and generating data, and dynamically interacting with other agents, bots, and humans. And despite the early days of AI adoption in most businesses, today, more than 50% of internet traffic is created by these autonomous workers.¹

A mind-blowing statistic like that reveals a dramatic disruption in business operating models and the demands they place on network infrastructure. Success will depend on the ability to move huge amounts of data from anywhere to anywhere, in real time, while delivering great customer experience at a reasonable cost. For your network to keep up, it must be big enough, fast enough, intelligent enough, and secure enough.

Consider that chip investment - the asset with the shortest economic half-life on earth. If it sits idle, money is wasted. Still, it's common practice in procurement departments to constrain bandwidth for cost control. Doing that also constrains data flow, starving precious GPUs. That's like filling your F1 race car with low octane fuel.

Bandwidth matters for customer experience, too. Employees, customers, partners - they don't like to wait. How quickly an AI system responds with insights, measured by time to first token (TTFT), is a key metric. Less bandwidth means longer time to insight...and gratification.

¹ The Definitive Source: Imperva/Thales Bad Bot Report (2025 Edition)

In networking, bigger is better. But who can afford enough bandwidth to support the near-infinite combinations of connections that AI agents demand?

No one at scale. At least not in legacy network architectures where connections are predetermined, point-to-point, and static, and bandwidth allocation is fixed.

To support the brave new world of AI, networks need to be completely adaptable, programmable, and consumption-based, just like cloud. Fire it up, shut it down. Make sure it's there, wherever and whenever you need it. Pay only for what you use.

Your network should be able to adapt to your business while being resilient to the environment around it. Hurricane heading towards that data center? Move the workload. Energy spot prices spiking? Move the workload. AI model needs data from a different cloud? Move it!

You can't do any of that on yesterday's networks.

So, make sure your network supports the future you're building. Your business depends on it. And so does the competitiveness of our country.

A handwritten signature in black ink that reads "Kate Johnson". The signature is fluid and cursive, with a long horizontal stroke extending to the right from the end of the name.

Kate Johnson
CEO, Lumen Technologies
lumen.com/isyournetworkready

