

---

## Connectivity: The Heartbeat and Nervous System of the AI-Era

Mallik Tatipamula and Vint Cerf

### **Introduction**

Artificial Intelligence (AI) may dominate today's headlines, but it does not stand on its own. Models and compute often capture the spotlight, yet they are only part of the story. AI operates within a much larger system. And in that system, *connectivity* serves as the heartbeat that keeps everything moving and the nervous system that allows it all to work together.

Here is the way to think about it: *AI needs data for training. Data relies on compute for processing. Compute in turn requires connectivity for data exchange and data transfer.* When all three work together, AI reaches its full potential. Connectivity is more than a utility. In the AI-era, it is the lifeblood that sustains intelligence and the wiring that allows it to function seamlessly.

### **Looking Back: Networks as the Circulatory System**

The history of the Internet underscores this point. Born as a research experiment in the 1970s, and 80s, early ARPANET links carried simple packets that enabled collaboration across labs. Data was scarce, compute rudimentary, and networks were narrowband, yet those early flows proved a new idea: intelligence could be shared across distance.

The 1990s brought the World Wide Web (WWW), and suddenly data was everywhere” with webpages, e-commerce transactions, streaming media. Search engines and recommendation systems flourished, but only because *connectivity* moved bits from servers to compute clusters. Without Wide Area Networking (WAN), information would have remained trapped in silos.

The 2000s layered on two revolutions: cloud computing and the smartphones. Cloud services delivered elastic compute, while smartphones put connectivity in billions of pockets. Social media, ride-sharing and mobile payments were not just software innovations; they were network-driven, requiring real-time coordination between devices, users and clouds.

In the 2010s, deep learning brought GPUs and massive models. But GPUs alone did not make AI practical. What elevated AI from an experiment to a global utility were networks: high-speed interconnects moving training data across distributed clusters, and mobile broadband delivering AI-powered services to billions of devices. The design clusters of connectivity: throughput, reach, latency, cost per bit, and reliability, were what allowed the organism to grow.

## **From Circulation to Coordination**

Connectivity's role has expanded from circulation (moving bits) to coordination (enabling reflexes).

The Internet of Things (IoT) in the 2010s highlighted this shift. Billions of sensors became the “nerve endings” of the digital world, streaming telemetry from homes, factories, vehicles and cities. Yet most IoT devices were passive: they sensed and reported but did not act in real-time. This exposed a gap. *Connectivity* was not just circulation but also neural wiring, essential for interpretation and coordinated response.

Emerging concepts such as the Internet of Senses aim to close this loop. By fusing sensing and communications through ISAC (Integrated Sensing and Communications), networks can become context-aware fabrics, transmitting what they detect in real-time. Multisensory technologies: haptics, digital olfaction, even brain-computer interfaces (BCI), extend communication beyond sight and sound, turning the Internet into a medium of experience rather than just information.

But perception without reasoning is incomplete. Here, AI agents emerge as the next step. Unlike IoT endpoints, agents are not passive. They perceive, reason and act. Digital agents such as copilots, workflow orchestrators, trading algorithms, live entirely in software. Physical AI agents like autonomous vehicles, drones, industrial robots, bring intelligence into the physical world. Both require connectivity not just as a circulatory system but also as a nervous system: wiring cognition across devices, edge nodes and clouds.

## **Why Connectivity Defines AI's Future**

Today's AI ecosystem demonstrates this dependency vividly.

At one extreme, foundation models contain trillions of parameters and run across thousands of accelerators in globally distributed data centers. Training such models requires high-bandwidth, low-latency interconnects like InfiniBand, Ethernet with RDMA, or emerging optical fabrics. Without these, multi-week training runs would be impossible.

At the other extreme, edge devices like smartphones, industrial sensors, medical wearables, now carry powerful NPUs and GPUs for on-device inference. Apple's Neural Engine, Qualcomm's AI Engine and Google's Tensor Processing Units (TPUs) in phones enable AI agents to run locally. But their usefulness depends on staying in sync with the cloud and peers through reliable connectivity.

What ties these extremes together is *connectivity* as fabric, spanning the system end-to-end:

- **Within data center:** Ultra-fast interconnects bind GPUs, TPUs, and accelerators for distributed model training

- **Across regions and continents:** High-capacity optical backbones move vast datasets and inference outputs globally.
- **At the edge:** Wired and wireless access networks bring intelligence to people, machines and environments
- **Beyond terrestrial limits:** Satellites and high-altitude platforms extend reach to underserved or remote regions.

This layered fabric ensures that data flows to compute when needed, compute delivers insights back in time, and intelligence emerges as a system rather than isolated silos. Without connectivity, compute is stranded. With *connectivity*, intelligence becomes collective: distributed across clouds, edge, and devices worldwide.

## **Technical Challenges Ahead**

If *connectivity* is the heartbeat and nervous system of the AI-era, then making it work at scale presents six major technical challenges.

1. **Ultra-low latency and determinism:** AI tasks such as autonomous driving, robotic surgery and industrial automation require sub-millisecond responsiveness with predictable guarantees. While 5G URLLC is a first step, AI-native networking must integrate sensing, scheduling and compute coordination far more tightly to ensure end-to-end determinism and real-time decision making.
2. **Bandwidth, fabrics and data movement:** Training trillion-parameter models produces exabytes of traffic, and today's Ethernet based interconnects and memory hierarchies cannot keep pace. Accelerators scale faster than I/O, leaving compute cycles stalled waiting for data. Breaking this bottleneck will require co-packaged optics, silicon photonics, rack-scale integration, and memory disaggregation (e.g., CXL) to deliver multi-terabit-per-second throughput per node and move data as efficiently as it is processed.
3. **Resilience and Security:** As AI workloads become critical infrastructure, connectivity fabrics must ensure fault tolerance and adversarial robustness. Multipath routing and self-healing meshes provide continuity under failure, while zero-trust models and AI-driven anomaly detection secure operations across cloud, edge and devices. Supply-chain integrity and quantum-resistant cryptography will be essential to sustain trust at global scale.
4. **Energy Efficiency:** From hyperscale data centers to radio access networks, connectivity is energy intensive. AI-Native networks must be designed with energy proportionality and sustainability in mind, ensuring performance without compromising sustainability goals.
5. **Orchestration and Interoperability:** Just as TCP/IP created a common foundation for the internet, the AI era offers a chance to establish open protocols for agent identity, inter-agent communication, and workload orchestration. Today's orchestration tools like Kubernetes for Cloud, MANO for NFV, O-RAN RIC for RAN, operate in silos. Moving forward, AI-Native

systems can unity these into an end-to-end framework that spans cloud, edge, and devices, ensuring seamless interoperability and preventing fragmentation.

6. **Trustworthy AI integration:** As networks themselves become AI-Native, ensuring the reliability, fairness, and explainability of AI-driven decision is paramount. From spectrum allocation to closed-loop control, bias or opaque inference could undermine trust. Embedding verification, validation and continuous monitoring of AI models into network operations will be critical.

## Lessons from History

History shows that breakthroughs in compute alone do not unlock progress: It is *connectivity* that turns isolated advanced into global transformations. The supercomputers of the 1990s were powerful but niche, and only when they were networked through the Internet, did the intelligence begin to scale across the world. The smartphone's true success also came not from its hardware alone, but from the power of always-on connectivity that enabled entire ecosystems of applications and services.

AI stands at a similar moment today. Compute will continue to advance, but its full impact will only be realized when paired with robust, open and ubiquitous *connectivity*. With this foundation, AI can grow into a planetary-scale utility: resilient, inclusive and transformative for society.

## Closing the Loop

When you step back, the patterns is striking:

- Data without compute is meaningless
- Compute without *connectivity* is stranded
- AI without both is nothing more than an idea

*Connectivity* has been there from the beginning: carrying packets, enabling mobility, linking machines. In the AI-era, it ensures that intelligence flows freely rather than remaining locked in silos. It synchronizes training across data centers, distributes inference to the edge, and coordinate agents acting in the physical world.

In summary, *connectivity is not just the foundation of AI. It is the pulse that keeps the system alive and the nervous system that makes it intelligent.*