

As Cloud AI spend passes \$800 billion, Future Grid turns idle customer devices into enterprise compute

Future Grid enables enterprises to run AI workloads on the idle devices their own customers already own, cutting cloud costs and paying users for contributing.

LONDON, UNITED KINGDOM, July 13, 2026 /EINPresswire.com/ -- The world's nine largest cloud providers are on track to spend close to \$830 billion in capital expenditure in 2026 (TrendForce), and the four largest US technology companies alone are guiding to roughly \$725 billion, most of it on AI infrastructure (Tom's Hardware). The strain is showing up at the socket: the International Energy Agency projects electricity demand from data centres will more than double to around 945 TWh by 2030, roughly Japan's entire electricity consumption today, with demand from AI-optimised data centres set to more than quadruple. In Ireland, data centres already consume 22 percent of all metered electricity.



[Future Grid](#)'s insight is that most of that demand does not need a warehouse. Inference, the everyday work of running models for users, accounts for an estimated 80 to 90 percent of AI's energy footprint (UNU-INWEH), and unlike training, it distributes. It can spread across ordinary machines, shift in time, and follow cleaner power when coordination is intelligent.

"Enterprises are paying twice for capacity they touch every day," said Jonathan MacDonald, founder and CEO of Future Grid. "They rent cloud GPUs at record prices while the phones, tablets and laptops of their own customers sit idle for most of the day. We let a company activate the compute fleet it already has a relationship with – and share the savings with the people who make it possible. The status quo pours compute into warehouses that ordinary people pay for but never see. We think the grid should give back more than it takes."



If your inference costs are climbing and your customers already open your app every day, the capacity you need is already in their pockets."

Jonathan MacDonald

The supply is not theoretical. Consumer networks such as Acurast already report more than 250,000 smartphone processors running production inference workloads in secure environments across 175 countries, from a single consumer app. A large enterprise with tens of millions of users can exceed that supply from its own customer base alone, provided participation is designed with clear value on both sides.

How Future Grid works:

- An enterprise creates a private network. Customers opt in inside the company's own app. A lightweight node runs on surplus CPU, throttled so the device stays responsive for its owner.
- The company routes its own inference and batchable AI tasks to those devices first. Work stays inside a private network by default: your nodes, your jobs, your brand.
- Work is verified automatically, with auditable proof of execution, and the company is billed against the hyperscaler cost it displaces, a fraction of the cloud bill.
- Users are rewarded for the surplus they choose to share, through bill credit, loyalty points or direct payment, funded from the savings the company would otherwise have sent to a cloud provider.

The environmental case is built in. Idle devices run on power already being consumed rather than new datacentre load, and because the work runs on hardware that already exists, it avoids the embodied carbon of manufacturing new servers. Researchers at Google and UC San Diego have shown that a modern smartphone processor core can match or beat a datacentre server core, and that roughly half of a phone's manufacturing emissions are locked in its motherboard alone (Google Research, 2026). Future Grid can also preferentially schedule work to cleaner grid windows, an approach shown in published research on carbon-aware scheduling (CarbonFlex) to cut emissions by more than half against a carbon-agnostic baseline. Each job carries a measured carbon record enterprises can use in ESG and sustainability reporting.

Future Grid is now seeking design partners: enterprises with an app-centric customer base and a rising AI compute bill, across energy, transport, telecommunications, banking, retail and media. Partners will embed the SDK, measure opt-in, and document savings against a hyperscaler baseline in a bounded pilot, co-shaping reward design and success metrics along the way.

"If your inference costs are climbing and your customers already open your app every day, the capacity you need is already in their pockets," MacDonald added.

About Future Grid

Future Grid is a technology company building the coordination layer that lets enterprises run AI workloads on the idle devices of their own opted-in customers. Its patent-pending platform provides the SDK, orchestration, verification, rewards and settlement to take an enterprise from pilot to production network, while privacy and control stay with the device owner. Future Grid is built on Celestial (<https://celestial.global>), which handles identity, verification and settlement as work moves between company systems and user devices.

For all enquiries: <https://thefuturegrid.org/>

All statistics cited in this release are drawn from published sources listed at <https://thefuturegrid.org/sources>

Jonathan MacDonald
Future Grid
info@thefuturegrid.org

This press release can be viewed online at: <https://www.einpresswire.com/article/926118754>

EIN Presswire's priority is source transparency. We do not allow opaque clients, and our editors try to be careful about weeding out false and misleading content. As a user, if you see something we have missed, please do bring it to our attention. Your help is welcome. EIN Presswire, Everyone's Internet News Presswire™, tries to define some of the boundaries that are reasonable in today's world. Please see our Editorial Guidelines for more information.

© 1995-2026 Newsmatics Inc. All Right Reserved.