From Idle Compute to Useful Compute

Incentive Design for Distributed Scientific Computation

Patent Pending
UKIPO Filing No: **GB2506299.3**
Filed: **27 April 2025**

Author: Che-Hwon Bae
Website: www.baemax.co.uk
Version: V1 — Revised Working Paper
Last Updated: January 2026

Abstract

Demand for computation in artificial intelligence, scientific research, and medical simulation is rising rapidly, while access to affordable, reliable compute remains constrained. At the same time, vast quantities of CPU and GPU capacity sit idle across homes, offices, and underutilised infrastructure. Current approaches to scaling compute rely heavily on the construction of large, centralised data centres, concentrating energy demand, capital expenditure, and environmental impact.

This paper argues that the primary barrier to activating latent compute is not technical feasibility, but incentive design and execution structure. We propose a framework for *useful compute*: a decentralised, incentive-aligned market that rewards computation based on real-world utility, reliability, and quality rather than raw throughput. The framework combines bilateral reputation, quality-weighted price discovery, persistent machine identity, and federated coordination pools to enable scalable, distributed execution of scientific and AI workloads. By aligning incentives with usefulness and system contribution, the approach offers a complementary alternative to centralised infrastructure expansion while improving resilience, environmental scalability, and access to compute.

Table of Contents

---

1. Introduction

Computation has become a non-discretionary input to modern science, medicine, and economic growth. Progress in artificial intelligence, drug discovery, climate modelling, and materials science increasingly depends on access to large-scale compute. Yet compute availability remains expensive, unevenly distributed, and concentrated among a small number of hyperscale providers.

Alongside this constraint exists a paradox: billions of CPUs and GPUs worldwide remain idle for substantial portions of time. Personal devices, office workstations, laboratory machines, and lightly loaded servers collectively represent a vast reservoir of latent compute. Despite this abundance, there is no global mechanism that activates this capacity for meaningful scientific or AI work.

This paper argues that the limitation is not hardware, but the absence of a market structure that prices *usefulness*, reliability, and coordination.

---

2. The Limits of Centralised Scaling

The dominant response to rising compute demand has been the construction of increasingly large data centres. While effective at scale, this model introduces systemic constraints:

- **Energy concentration:** Large facilities impose step-function increases in local power demand, stressing grids and requiring significant upstream investment.
- **Capital intensity:** Data centres require long lead times, regulatory approval, and substantial upfront expenditure.
- **Environmental and political friction:** Concentrated infrastructure attracts opposition and can distribute costs regressively at the local level.
- **Brittle scaling:** Capacity is added in discrete increments rather than smoothly.

These characteristics mirror execution failures observed in other large-scale transitions, where infrastructure expansion outruns system absorption capacity. The challenge is not compute itself, but how compute is activated and scaled.

---

3. Latent Compute as a Systems Resource

Latent compute differs fundamentally from newly built infrastructure:

1. **It already exists** — no fabrication, land acquisition, or major grid upgrades are required.
2. **It is geographically dispersed** — energy draw is distributed rather than concentrated.
3. **It is heterogeneous** — machines vary in capability, reliability, and suitability.

Many environments already operate machines that remain powered for extended periods, such as office PCs, laboratory workstations, and on-premise servers. Useful compute allows a configurable portion of idle resources to be contributed opportunistically, subject to local policy constraints on CPU, GPU, memory, priority, and time windows. Contributions occur only when machines are otherwise underutilised, preserving primary use while enabling incremental revenue generation without behavioural change or dedicated hardware investment.

Activating this resource is therefore a market design and execution problem rather than a technical one.

---

4. Why Existing Models Fall Short

4.1 Hyperscale Cloud

Centralised cloud platforms abstract hardware heterogeneity but internalise coordination while externalising environmental and capital constraints.

4.2 Decentralised GPU Rental Markets

Decentralised marketplaces (e.g. Vast.ai) demonstrate that distributed compute can operate at scale, but largely treat compute as a homogeneous commodity priced per GPU-hour. Reliability, coordination, and execution quality are left to the requester, creating friction and adverse selection.

4.3 Crypto Mining

Cryptocurrency mining successfully mobilises distributed compute but optimises for a narrow objective function unrelated to real-world utility, consuming energy without producing socially useful outcomes.

None of these models price *usefulness*, *reliability*, or *coordination* explicitly.

---

5. Comparative Positioning

Dimension	Hyperscale Cloud	GPU Rental Markets (e.g. Vast.ai)	Crypto Mining	Useful Compute
Ownership	Centralised	Distributed	Distributed	Distributed
Pricing Basis	Usage	GPU-hour	Hash rate	**Utility & quality-weighted**
Coordination	Centralised	Minimal	None	**Federated pools**
Quality Signalling	Internal	Weak	None	**Bilateral reputation**
Incentive Alignment	Cost	Capacity	Token reward	**Outcome & reliability**
Energy Profile	Concentrated	Mixed	High & wasteful	**Dispersed & incremental**
Scaling Behaviour	Step-function	Limited	Poor	**Smooth / absorptive**

---

6. Defining Useful Compute

Useful compute is computation whose value derives from contribution to real-world outcomes rather than raw throughput. Relevant dimensions include:

- execution reliability,
- availability and uptime,
- reproducibility of results,
- task suitability,
- and alignment with project objectives.

A system that rewards useful compute must embed quality directly into price discovery.

---

7. Bilateral Reputation and Quality-Weighted Pricing

The framework operates as a bilateral matching market between compute contributors (“workers”) and task requesters (“employers”).

Contributor evaluation includes execution reliability, task completion, uptime consistency, and reproducibility.
Requester evaluation includes task clarity, payment reliability, fairness, and operational behaviour.

Price emerges endogenously:

- lower-quality requesters must offer higher compensation to attract reliable compute,
- lower-quality contributors receive lower compensation for equivalent tasks.

Rather than exclusion, quality differences are internalised through price.

---

8. Persistent Machine Identity

To prevent reputation reset and ensure continuity, compute contributions are associated with persistent machine identities, analogous to enterprise software licensing. Reputation follows execution capability rather than user accounts alone.

Hardware changes are tracked, allowing the system to:

- distinguish genuine upgrades from opportunistic resets,
- model degradation or improvement over time,
- maintain auditability without invasive monitoring.

---

9. Federated Pools and Delegated Coordination

Coordination overhead is a major barrier to decentralised execution. The framework therefore supports federated compute pools.

- Any individual or organisation may operate a pool.
- Pools aggregate contributors under shared standards.
- Membership is voluntary and competitive.

Pool coordinators act as execution managers, handling task distribution, verification, retries, and quality enforcement. In return, coordinators receive a transparent share of compute rewards or project outcomes.

From the requester's perspective, interacting with pools rather than individual machines dramatically reduces operational burden while preserving decentralisation.

---

10. Security and Data Protection Considerations

Security challenges are not novel; they are issues of execution and policy enforcement rather than cryptographic invention.

From the requester's perspective:

- workloads run in sandboxed containers or virtual machines,
- data is encrypted in transit and at rest,
- optional confidential compute mechanisms can be used,
- verification relies on redundancy and reproducibility.

From the contributor's perspective:

- raw data is not exposed,
- resource usage is strictly capped,
- workloads are transparently classified,
- contributors may opt out of task categories.

Trust is enforced structurally rather than socially.

Many of the technical building blocks required for distributed execution — including containerisation, distributed training, and federated learning — are well-established in the literature and industry. This work focuses on incentive structure and execution coordination rather than algorithmic innovation.

---

11. Environmental and Execution Implications

By activating existing hardware rather than constructing new facilities, useful compute:

- avoids concentrated energy demand,
- scales incrementally,
- reduces political and environmental friction.

This mirrors execution lessons from other infrastructure transitions, where distributed contribution and incentive alignment outperform monolithic build-outs.

---

12. Future Extensibility

Modern consumer and embedded devices increasingly contain idle general-purpose compute capacity. While practical constraints currently limit participation, these systems illustrate the scale of latent computation embedded in everyday infrastructure. The useful compute framework is designed to be extensible as isolation, policy control, and verification mechanisms mature.

---

13. Conclusion

The world does not lack compute; it lacks mechanisms to activate it usefully. Centralised scaling alone risks repeating execution failures seen in other large-scale transitions, concentrating cost and fragility.

By treating compute as a market shaped by incentives, reputation, and coordination, useful compute offers a complementary path — one that scales smoothly, rewards real contribution, and aligns execution with societal value.

Status and Disclosure

This paper describes a conceptual framework protected under a provisional patent filing with the UK Intellectual Property Office (Filing No. GB2506299.3). The discussion focuses on market structure and execution principles rather than implementation details.