THE WORLD IS THE MAINFRAME

A vision for shared global computing and storage

Concept developed circa 2006 | Discussion document, October 2026

Core proposition

Why build ever more dedicated computing and storage infrastructure before making better use of the resources already distributed across the world? Computers owned by individuals, businesses and institutions could contribute their unused storage and processing capacity to a shared network. Participants retain control of their own machines and benefit from the network when they need resources themselves.

“The world is the mainframe.”

1. The starting point: resilience through distribution

The original inspiration was the resilience of packet-switched networking: information can travel by different routes when a connection fails. The proposed extension is to make storage and computation similarly independent of any single physical machine.

2. Files as distributed packets

A file is divided into many small, identifiable packets or chunks. Each has an assigned position in the original file. These packets can be stored across numerous independent hard drives rather than relying on one server to hold the whole file. When the file is requested, the network retrieves the necessary packets and reconstructs them in order.

A distributed header or index describes the file structure, packet identities, order and locations. Copies of this index are held in multiple places so the information needed to reconstruct a file is not itself a single point of failure. Individual packets are also replicated across different machines. If one node disconnects, alternative copies can be retrieved; the network can then replenish redundancy.

When file contents change, only the affected packets need updating, rather than transmitting and storing the entire file again. In a practical implementation, versioning and consistency rules would be required to prevent mismatched packets or indexes.

Concept and design by Christopher Crawford

3. Spare storage as a shared resource

People and organisations could rent out unused hard-drive capacity, much as owners of solar panels can contribute surplus generation to a wider energy system. They could withdraw capacity when needed, with the network designed to withstand such departures through replication and redistribution. Participation would be voluntary and compensated.

4. The world as a distributed processor

The same principle applies to computing. A large task can, where technically possible, be divided into smaller units and assigned to many machines. A fast processor might complete thousands of units while an older, slower computer completes only a few. Each contribution remains useful. Results are checked and assembled into the final output.

As working days end in one part of the world, spare computing capacity could become available there. As those machines are needed again, suitable jobs could be checkpointed, reassigned or resumed elsewhere. The network would continuously match work to available machines rather than depend on a fixed fleet of servers.

5. Mutual benefit, not charity

The economic principle is simple: “If you share, you all gain.” A company can earn value from idle infrastructure and gain access to extra capacity during periods of high demand. A participant can be both a provider and a customer. Possible models include payments for verified capacity and completed work, or resource credits that can be exchanged for future usage. The original discussion establishes the principle of mutual benefit; it does not prescribe a final payment model.

6. Why this matters

- Better utilisation of computers, processors and storage already in existence.
- Potentially lower barriers to accessing large-scale computing resources.
- Less dependence on individual storage servers or computing facilities.
- A new revenue stream for owners of underused equipment.
- Potential to extend the useful life of older machines and reduce demand for new hardware, subject to energy-efficiency analysis.

7. Practical questions to solve

Reliability: How many packet copies are required, and how quickly must lost replicas be replaced?

Security and privacy: How are data encrypted, keys controlled and untrusted machines isolated?

Integrity: How does the network verify that stored packets and completed calculations are genuine?

Coordination: How are indexes, jobs, checkpoints and versions maintained without creating a central point of failure?

Performance: Which tasks divide efficiently, and when do latency and bandwidth outweigh the benefit of distribution?

Economics and energy: How are contributors compensated after accounting for electricity, network traffic, wear and administration?

8. Relationship to existing technology

Distributed storage, volunteer computing, peer-to-peer networks and resource marketplaces each demonstrate aspects of this vision. Examples include BOINC-style distributed computation and later decentralised storage services. These parallels do not establish that the original proposal was identical to any particular system, nor that every part is technically or commercially solved.

9. The vision in one paragraph

A global network in which people and organisations voluntarily contribute spare storage and computing power; files are split into replicated, indexed packets; suitable processing tasks are broken into units that machines of different capabilities can complete; work moves as machines become busy or unavailable; and contributors share in the value created. The aim is not to own more infrastructure, but to use the world's existing infrastructure more intelligently, resiliently and fairly.

Next questions

A fuller specification could define the packet/index format, replication policy, job scheduling, verification, encryption, participant rewards and a small working prototype. This document records the concept as described in conversation; it is not a patent assessment, a proof of priority or a claim that a general-purpose global network has already been demonstrated.