Bringing It Home

The case for the sovereign edge · home servers, honest ledgers, and a distributed internet built for ordinary people

By Will Poynter · Porter Valley Technology Company · v1.0

Abstract

For thirty years we have rented our digital lives. Our photos, messages, records, and creative work sit in other people's data centres, where we trade ownership for convenience and, in doing so, enrich whoever holds the data. This paper argues that the next era of the internet moves storage and computation into the home, to a home server as ordinary as a boiler, and that doing so is not only a matter of privacy but of national capability: more sovereign compute, more secure data, and a route to net-zero that building more hyperscale data centres cannot offer. The barrier has never been desirability, but usability. We have never made the case for why an ordinary household would want a server, nor built one accessible enough for them to adopt.

---

1. We are tenants, not owners

The bargain of the cloud looks fair: hand your data to a data centre and receive reliability and reach in return. Your photos are safe and available from any device, backed up better than you could manage yourself. In return, we accept that these companies are able to make money from our data, either through direct selling or by deriving new intellectual property. This agreement has some parallels with leaving our money at a bank to keep it secure, and while they keep the money they can use it to make more money. However, banks are mandated to tell us how much of that new money they will share with us via interest, the data and intelligence economics of the platforms have not been subject to the same regulation.

The issue does not stop at the software layer, but goes down to the physical substrate. Each layer of the system can be charged to the next. Once users are comfortable with the platform and servers, it becomes easier to sell other layers of product whether that is SecOps in the case of Amazon Web Services, or the Marketplace on Facebook. The more layers an organisation owns, the more revenue it can derive. Owning the physical layer, the first layer, is rewarded with the most control and the highest efficiency.[1]

The result has been an oligarchy. A handful of men hold the records of billions of lives, and that concentration now distorts the very democracies that host it. Governments have noticed and the UK government has responded by marshalling billions in investment toward data centres.[2][3] While this may address some of the UK's goals, it does not address the issue of wealth extraction across borders.

2. The edge is the answer

To remain competitive globally, which is vital for powering the NHS and defence, we need to increase the compute capacity of the UK. To improve the security and safety of our data, we need the data stored on drives under UK control. Sovereign data centres is the natural answer, and may be part of the solution, but it also works against several other of the government's goals around not building on the green-belt and achieving net-zero.[4]

By locating more of the compute and storage truly at the edge, we are able to still achieve our aims above of increasing compute and sovereign data, while also achieving several others. Households and individuals have the ability to store their online presence where they can control it and choose if and how it is monetised. For example, loaning their anonymised wearable data to a researcher's study. By relocating control, we make the Internet more resistant to capture.

Inference is computationally more expensive than other tasks typically performed by an end user on a daily basis and in turn the most expensive in electricity and heat dissipation.[1][5] While training is also expensive, most people will not expect to train AI, while other computation tasks like sending an email, spell checking or adding a filter to a video are all trivial by comparison on current hardware. Electricity used by a server is almost entirely ejected as heat, which in UK households can be put to good use heating the home or heating water. Integration with heat pump systems would be a natural fit to have the most efficient climate control.

Electricity can be used most efficiently when its generation is co-located with its use due to avoiding transmission loss.[6] Fitting solar panels and batteries to homes can be used for localised, green and cost efficient generation most of the year.

We still need some new centralised facilities in the UK, but only to do the most challenging research and development work. Day-to-day use of the expanding digital market requires trust for its customers. Untangling the two issues is key to solving them both efficiently. Most of the inference for mundane tasks and personal data go to the edge of the Internet, people's homes and offices. Developing new models should be directed towards supercomputers while industry, hospitals, universities, and governments require mid-tier capacity to host the multi-trillion parameter models that will perform advanced automation.

This is not only a technical preference; it is a trust strategy designed to unblock AI-enabled growth. A government trying to deploy ambitious technology, digital identity, health data sharing, AI in public services, needs public trust to spend, and that trust is far easier to earn when a

citizen's data sits in their home rather than a distant vault. Sovereignty at the edge is how you rebuild the trust that centralisation spent.

3. Distributed technologies need a home

Moving computation to the edge does more than hand a household control of its own data; it lays the foundation for a wider class of tools known as distributed technologies. Their appeal is straightforward: a network with no single owner is harder for any one party to capture, control, or shut down, and leaves no single gatekeeper positioned to charge everyone who uses it. Distributed storage, distributed identity, and distributed ledgers all belong to this family. Many aspects of the underlying technology have become entangled in the public relations disaster of cryptocurrencies.

Distributed systems require devices, typically home computers or servers, to be spread out across many different locations. Functioning at their best requires a large number of independent computers that are switched on, online, and dispersed widely enough that no small group controls them. Today, blockchain servers are typically run by tech-savvy people or companies in data centres and mining centres. The consolidation of the nodes under one organisation’s control undermines the purpose of distributed computing from a trust perspective.

Equally important, the distribution of the governance of any distributed technology will not be tolerated globally by a local-first philosophy. Whether a distributed technology is: controlling when public records are accessible automatically, determining who has break-glass authority to enter a data vault for an investigation, or allowing the use of AI generated content, communities have differing opinions that are much easier to solve regionally. If we apply these two guidelines, we want lots of different owners of servers spread out in regional clusters. Some networks could be federated depending on their function.

4. Resilience

From discussions and interviews with non-technical users, the most common worry is backup: what happens if their server is lost permanently? Today the cloud answers this by holding several copies of each file across distant data centres, so a burglary or flood at home loses nothing.

A home-server network has to solve the same problem with servers that are less reliable and not always friendly. The approach is to split each file into encrypted fragments and spread them across many homes, using an erasure coding technique, such as Reed-Solomon[7]. This adds redundancy: the original file can be rebuilt from a subset of the fragments, so the data survives even when several host servers are offline. Encryption does a separate job; without the owner's

key, a fragment is meaningless to whoever is storing it. Together, the two give you availability and secrecy: your data is there when you need it, and no host can read it.

On top of this sits a governance layer that enforces the rules. Some of those rules are purely technical: how many fragments must exist, how long one can be missing before the network rebuilds it. These questions interest computer scientists and almost no one else.

But a second set of rules is not technical at all. Who may access a file? Can an authority demand it, and what must they prove before they can? Those are political questions, and they have different answers in different communities. The key design decision is to keep the two apart. Settle the engineering issues separately and make the access rules visible, auditable, and locally governed. Most systems tangle them together, which is precisely why people don't trust them. Separating them is how trust is earned.

5. What we get

Owning a server in the digital economy is the equivalent of holding land in the real world. A citizen with land can live on it, farm it, rent it out, build a mill or a factory. A household with a server can do the digital equivalent: host their own data, decide whether and how it is monetised, and rent spare capacity to others — a new marketplace running between homes rather than through a distant landlord.

There will always be a need for data centres. But the household would no longer have to place so much trust in them, because the things that matter most would already be home.

References

[1] L. A. Barroso, J. Clidaras, and U. Hölzle, *The Datacenter as a Computer*. Morgan & Claypool, 2013.

[2] “A decisive shift to power British AI: new £1.1 billion plan to back chip firms, boost computing power and skills for the AI revolution - GOV.UK”, *UK Government*, 2026. [Online]. Available: https://www.gov.uk/government/news/a-decisive-shift-to-power-british-ai-new-11-billion-plan-to-back-chip-firms-boost-computing-power-and-skills-for-the-ai-revolution. [Accessed 18 Jun 2026]

[3] “ Building Britain's AI Future | Sovereign AI Mission”, *UK Government*, 2026. [Online]. Available: https://www.sovereignai.gov.uk/mission. [Accessed: 18 Jun. 2026]

[4] N. Burnett and I. Stewart, "CBP-9888," House of Commons Library, Jun. 2026. [Online]. Available: https://researchbriefings.files.parliament.uk/documents/CBP-9888/CBP-9888.pdf. [Accessed: 18 Jun. 2026]

[5] F. Oviedo, F. Kazhamiaka, E. Choukse, A. Kim, A. Luers, M. Nakagawa, R. Bianchini, and J. M. Lavista Ferres, "Energy Use of AI Inference: Efficiency Pathways and Test-Time Compute," Microsoft, arXiv:2509.20241, Sep. 2025. [Online]. Available: https://arxiv.org/pdf/2509.20241. [Accessed: 18 Jun. 2026]

[6] UK Power Networks, "Distribution Network Energy Losses," UK Power Networks, 2024. [Online]. Available: https://www.ukpowernetworks.co.uk/distribution-network-energy-losses. [Accessed: 18 Jun. 2026]

[7] I. S. Reed and G. Solomon, "Polynomial codes over certain finite fields," SIAM J. Appl. Math., vol. 8, no. 2, pp. 300–304, 1960.