

The Other Answer

A Cube Commons response to The Technological Republic

Cube Commons Technical Report CC-TR-2026-005 Barton Nicholls · April 2026

I.

Alexander Karp and Nicholas Zamiska have written a bestseller that poses an important question and gives the wrong answer. *The Technological Republic* asks how democratic societies should organize the relationship between software infrastructure and institutional power in the age of artificial intelligence. Its answer: fuse Silicon Valley to the Pentagon, concentrate capability in state-aligned champion firms, and revive what it calls the “engineering elite” to build a new Manhattan Project. The book has been endorsed by Walter Isaacson, Eric Schmidt, Jamie Dimon, and George Will. *The Times of London* calls it “the AI manifesto inspiring Keir Starmer’s government.” It is absorbing the policy oxygen at exactly the moment when civic institutions are making decisions about AI infrastructure that will shape public life for decades.

We think it is the wrong answer. Not because the question is wrong — it is the right question — but because the book’s architectural prescription inverts four centuries of durable democratic practice, ignores forty years of Nobel-Prize-winning empirical political economy, and misreads the actual foundational vision of computing. Most consequentially, it is a prescription being written against a live counter-example that Karp and Zamiska do not engage with: between 2018 and 2025, a polycentric ecosystem of commons-governed software defeated serial attempts by well-resourced firms to enclose critical infrastructure. The state-aligned concentration model is not gathering empirical wind. It is losing, in real time, in exactly the infrastructure layer its proponents claim it alone can secure.

Cube Commons exists to build the other answer. This is a sketch of what that other answer is, and why we think it is the more durable democratic architecture.

II. What Karp and Zamiska actually argue

The book’s central move is to diagnose Silicon Valley as suffering from a patriotism deficit. The engineering class, in this telling, has retreated into photo-sharing apps and marketing algorithms when it should be building weapons systems, intelligence platforms, and civic

infrastructure for the United States and its allies. The remedy is a renewed compact between government and private technical elites, with Palantir offered — sometimes explicitly — as the existence proof that such a compact can work.

Beneath the patriotic vocabulary, the architectural claim is narrower: that durable democratic infrastructure emerges from concentration of AI and data capability inside state-aligned private firms. Karp is a skilled writer and does not say this in so many words. But the operational picture — Gotham, Foundry, Apollo, AIP; the £330 million NHS Federated Data Platform; the \$10 billion US Army enterprise agreement consolidating 75 contracts; the \$30 million ICE ImmigrationOS — is what the book is arguing for when it argues for anything concrete. The claim that the engineering elite must “rebuild its relationship with government” is, in practice, a claim that public institutions should outsource their core data and analytic infrastructure to a small number of vendors whose products are proprietary, hosted, and whose internal “Ontology” layer functions as a permanent dependency.

There is an honest version of this argument and a dishonest one. The honest version is: concentrating capability in firms with strong internal engineering cultures produces better software than fragmented public procurement. This is sometimes true. The dishonest version is that such concentration is therefore the *democratic* architecture — that it preserves the institutions it hollows out. This is the version the book prefers, and it is the version we want to contest.

III. The commons that don't collapse

The deep assumption underneath *The Technological Republic* is one Karp and Zamiska never name, probably because they have absorbed it too thoroughly to see it. It is Garrett Hardin's 1968 frame: unmanaged commons must collapse; the remedy is private property or state command. Fifty years of policy discourse absorbed Hardin uncritically. Karp and Zamiska have inherited the conclusion — that distributed institutional arrangements are fragile and must be backstopped by concentrated power — without inheriting the knowledge that the foundational paper was wrong.

Elinor Ostrom spent her career demonstrating it was wrong. *Governing the Commons* (1990) walks through Swiss alpine grazing documented by charter since 1224, Japanese *iriai* forests, Valencia's Tribunal de las Aguas adjudicating irrigation disputes every Thursday since 1435, Philippine *zanjera* irrigation, Southern California groundwater basins. Real people, governing real shared resources, for real centuries. Her eight design principles — clearly defined boundaries, rules fit to local conditions, collective choice, accountable monitoring, graduated sanctions, low-cost conflict resolution, external recognition, and nested enterprises — are not a theoretical preference. They are an empirically derived architecture for durable institutional governance at scale, and they earned her the Nobel

Memorial Prize in Economics in 2009.

The concept Ostrom formalized is **polycentric governance**: many autonomous, rule-governed centers of decision-making, nested, coordinating through shared institutions, adapting faster than any single hierarchy can. This is not eccentric. It is, by now, the empirically best-supported model of durable institutional arrangements that political economy has. It is also a near-perfect description of the actually-functioning large-scale digital infrastructure we all rely on: the IETF, the W3C, the Apache Software Foundation, the Linux Foundation, Debian, the Internet itself. "We reject kings, presidents, and voting," said David Clark at IETF 24 in 1992. "We believe in rough consensus and running code." That is how the protocol layer that underwrites the entire software economy is governed. It works.

IV. Subsidiarity is four centuries old and is European constitutional law

The political principle that fits this architecture is subsidiarity. It predates modern computing by four hundred years. Johannes Althusius published *Politica Methodice Digesta* in 1603, treating political society as nested *consociations* — family, collegium, city, province, commonwealth — in which sovereignty resides in the people and is delegated upward only when the lower body cannot accomplish the task. Pius XI gave it canonical form in *Quadragesimo Anno* in 1931, explicitly against both Soviet statism and Mussolini's corporatist fusion: "it is an injustice and at the same time a grave evil and disturbance of right order to assign to a greater and higher association what lesser and subordinate organizations can do." Written in 1931, against exactly the state-firm fusion *The Technological Republic* now proposes as democratic salvation.

Subsidiarity is not a theological preference. It is justiciable European Union constitutional law under Article 5(3) of the Treaty on European Union and the Lisbon Early Warning System. It has been claimed across the political spectrum: by Catholic social thought, by Burke's "little platoons," by Tocqueville's voluntary associations, by Nisbet's intermediate institutions, by MacIntyre's communitarianism, by Ostromian political economy, by the distributist left. The diagnosis that concentrated power hollows out the institutions it claims to serve — and that durable democratic life requires strong intermediate bodies holding their own capacities, including their own data and their own computation — is genuinely cross-partisan.

The Technological Republic never engages with any of this. It proceeds as though the political choice were only between Silicon Valley consumer frivolity on one hand and state-aligned champion-firm fusion on the other. There is a third option, and it is the older one.

V. The tools matter, and Palantir-class tools disable the

institutions they augment

Ivan Illich, in *Tools for Conviviality* (1973), gave us the vocabulary we need. Every tool, he argued, yields genuine benefit up to a first threshold and then, past a second threshold, disables the competence it ostensibly served. Past the second watershed, a tool achieves what he called **radical monopoly** — not one vendor beating another, but the foreclosure of the alternative form of practice entirely. “That motor traffic curtains the right to walk, not that more people drive Chevies than Fords, constitutes radical monopoly.”

Palantir-class platforms are paradigmatic second-watershed, radical-monopoly tools. A hospital system that places its data inside Foundry’s Ontology has not augmented its analytic competence; it has transferred a core civic capacity to a vendor whose architecture ensures the competence cannot be rebuilt internally. A police department running Gotham has not strengthened its institutional judgment; it has restructured judgment around the tool. The same applies at the national scale: the NHS Federated Data Platform is not a tool the NHS owns and operates. It is a reorganization of the NHS around a vendor’s ontology.

This is what Lessig meant in 1999 when he said “code is law.” Architecture is constitutional choice. When the code of a civic institution’s most important analytic infrastructure is proprietary, hosted, and built around vendor-specific data schemas, the institution has delegated a constitutional function to a commercial actor. The delegation is usually invisible in procurement documents and fully visible five years later, when the institution tries to leave.

VI. The actually-foundational vision of computing was distributed

The Technological Republic leans heavily on a historical claim — that Silicon Valley’s origin was Pentagon-funded and that its drift from state service is the deviation to be corrected. This is roughly a third of the truth. The other two-thirds: the actually-foundational intellectual vision of computing, the one the builders themselves articulated, was distributed augmentation, not centralized replacement.

Norbert Wiener, in *The Human Use of Human Beings* (1950), warned about exactly the consolidation Karp and Zamiska now advocate. In 1949, he wrote to Walter Reuther at the UAW offering to help organized labor prepare for automation, and refused to share his work with the military. J.C.R. Licklider’s 1960 paper “Man-Computer Symbiosis” and his 1968 follow-up with Robert Taylor, “The Computer as a Communication Device,” frame computing as a network of autonomous nodes augmenting human judgment — the direct ancestor of the Internet’s end-to-end principle. Douglas Engelbart’s 1962 framework “Augmenting Human Intellect” analyzed computing as a socio-technical ensemble with recursive

institutional bootstrapping at its center. Ted Nelson's Xanadu encoded two-way links, transclusion, and fragment-level provenance at the protocol layer in 1960. Stafford Beer's Project Cybersyn in Allende's Chile (1971–1973) was nested-recursive computational coordination respecting subsidiarity by design.

The distributed-augmentation tradition is not a minor dissenting strand. It is the actual intellectual foundation. The state-aligned concentration model Karp and Zamiska treat as the tradition is a late deviation. The book's historical frame works only if you forget the actual history.

VII. The live evidence: polycentric commons governance is winning

If this were only a theoretical dispute, reasonable people could disagree and move on. But the dispute is being settled in real time, in the exact infrastructure layer that matters, and the commons-governance side is winning.

Between 2018 and 2025, four well-resourced commercial open-source firms attempted to enclose commons they had helped build: MongoDB in October 2018, Elastic in January 2021, HashiCorp Terraform in August 2023, and Redis in March 2024. Each relicensed away from OSI-approved open-source terms in an attempt to prevent Amazon Web Services and other hyperscalers from reselling their work. Each triggered a polycentric response. AWS forked OpenSearch from Elasticsearch in April 2021, now governed by the Linux Foundation. The Terraform community forked OpenTofu, reaching general availability in January 2024 under Linux Foundation stewardship; OpenBao followed for Vault. The Redis community forked Valkey thirty days after the relicensing; within a year, 83% of large Redis users were on or testing Valkey.

Then the reversals. Elastic re-adopted AGPLv3 in August 2024. Redis re-adopted AGPLv3 in May 2025. A November 2024 study by Brockmeier and colleagues at CHAOSS and the OpenForum Academy found no evidence that the relicensing improved revenue trajectories for the firms that attempted it.

What happened here matters. A decentralized ecology of neutral hosts (the Linux Foundation), legal backstops (AGPLv3, a license Stallman and Moglen wrote in 2007 specifically for this scenario), enterprise procurement pressure, and fast-coalescing community forks beat some of the best-funded attempts at enclosure of the last decade. Polycentric commons governance defeated concentrated firm power in the infrastructure layer where it actually lives. There is no comparable contemporary record on the state-aligned concentration side. Karp and Zamiska cannot point to one, because it does not exist.

VIII. The Cube Commons thesis

The architecture Cube Commons builds rests on four commitments, each of which has prior art going back centuries but which, taken together, constitute a position distinct from both Silicon-Valley-as-usual and the Karp/Palantir fusion model.

First, institutions own their data and their computation. In the CUBEdesk architecture, the building's data lives in the building's own cloud account. The operator runs mission control in a separate account. Cube Commons, upstream, has no runtime custody. This is the Ostromian boundary principle and the Illichian conviviality principle made operational in software. It is also what Estonia's X-Road has done at nation-state scale since 1999: ministries build their own systems, coordinating through a peer-to-peer data-exchange layer, with no central hub.

Second, the code is open-source, and it stays open-source. We use copyleft licensing — AGPLv3 and its kin — because forty years of free-software practice has shown that durable commons require self-executing legal instruments, not good intentions. We are Public Benefit Corporation structured because commitments that depend on quarterly discretion are not commitments.

Third, coordination happens through open protocols, not proprietary ontologies. The Cube Data Plane specification is a protocol, not a platform. Other people can implement it. This is what the IETF tradition does. It is what the W3C does. It is what made the Internet what it became, as against the closed alternatives (Prodigy, AOL, Minitel, X.25) that looked stronger in 1993 and are now footnotes.

Fourth, the institutions coordinated this way are real, named, and accountable. Distributed institutional sovereignty is not crypto-secession. It is not exit to the cloud. It is not the Network State. The institutions are existing civic bodies — hospitals, housing authorities, school districts, state legislatures, building associations — operating their own infrastructure under their own governance, coordinating horizontally through open protocols and vertically through the subsidiarity-respecting public institutions they already sit within. This distinction matters, because the hard-power right and the crypto-exit right both claim “decentralization” while meaning almost opposite things. We mean neither.

IX. What we are actually asking for

The Technological Republic frames the choice as between shallow consumer software and patriotic weapons platforms. That framing is false and should be rejected. The real choice, the one civic institutions are actually making right now in procurement meetings and AI strategy documents, is between two architectures of durability.

One architecture locates democratic resilience in the state's alignment with a few

concentrated firms whose proprietary ontologies become the de facto constitutional layer of public administration. This is the architecture Karp and Zamiska propose. It is the architecture that, in April 2026, is absorbing most of the policy attention in Washington, London, and Brussels.

The other architecture locates democratic resilience in the nested, distributed, rule-governed institutions that have sustained complex coordination for centuries — updated with the legal instruments (copyleft), the architectural commitments (local-first, end-to-end, federated), and the governance practices (polycentric, subsidiarity-respecting, commons-based) that four decades of free-software work and forty years of Ostromian political economy have made available.

The second architecture has deeper roots, a longer track record, and is currently winning the empirical contest in the infrastructure layer where the contest actually happens. Its main weakness is that it has been less well articulated as a public thesis than the first. This paper is part of fixing that.

There is another answer to the question Karp and Zamiska are asking. It is older, sturdier, more democratic, and already building. Our job, and the job of the institutions we build with, is to make sure it is understood in time.

Cube Commons, Inc. is a Massachusetts Public Benefit Corporation building local-first, open-source multi-agent coordination infrastructure for civic institutions. This paper accompanies Technical Report CC-TR-2026-005 and is deposited on Zenodo under CC BY 4.0.