

A Reader's Primer

An addendum to "The Other Answer"

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

How to use this primer

"The Other Answer" moves fast through ideas with long histories. This addendum slows down and walks through each of them for readers who haven't encountered the sources before, or who have met them in passing but never in connection with one another. It is organized in three tiers.

Tier 1 — For the general reader. What the concept is, in plain language, with an example from ordinary life. If you read nothing else, read these. They are enough to follow the argument.

Tier 2 — For civic technologists and practitioners. How the concept maps onto software, infrastructure, and institutional design. If you work in civic tech, government digital services, or open-source, this tier is where the operational content lives.

Tier 3 — For policy readers and legislators. Where the concept sits in law, political economy, and the actual policy debates happening right now. If you are considering grants, legislation, or procurement, this is where the decisions get made.

The concepts are presented in the order they appear in the main paper. You can read the primer front to back, or jump to whichever term you stumbled on. Nothing here assumes you have read the main paper's footnotes, though references are given at the end.

1. The commons

Tier 1

A **commons** is a shared resource that a community governs together. The word comes from medieval English villages where the grazing land outside the hedges was held in common — every villager had the right to graze a fixed number of animals on it, and the village collectively decided the rules, enforced them, and adjusted them over time when conditions

changed. There were commons for pasture, for woodland, for fisheries, and for water.

When Americans hear “commons” today, most of us think of Boston Common or the town green in New England. That is the descendant. The deeper idea, though, is not “public space” but “shared resource governed by its users according to rules they make themselves.”

The important thing to understand is that commons are **not the same as open access**. Open access means anyone can take as much as they want, with no rules. A genuine commons has rules, boundaries, and enforcement — the rules are just made by the community of users rather than by a private owner or a distant government.

Tier 2

For civic technology, the commons framing matters because most of the internet’s foundational infrastructure — the Linux operating system, the Apache web server, the protocols like HTTP and TCP/IP that make the internet work — is governed this way. Nobody owns Linux. Nobody owns the internet protocols. They are maintained by communities of contributors under rules those communities wrote, and they have outperformed every proprietary alternative at the infrastructure layer for thirty years.

When we say Cube Commons is building “commons infrastructure for civic institutions,” we mean infrastructure that behaves the way Linux behaves: open, community-governed, not owned by any single firm, available for anyone to inspect, modify, or fork if the governing community goes bad.

Tier 3

Policy readers should know that the commons tradition is not fringe. Elinor Ostrom won the 2009 Nobel Memorial Prize in Economics for showing that commons are among the most durable institutional arrangements humans have built, sometimes lasting eight hundred years or more. Her eight design principles (see §3 below) are now taught in public-administration programs worldwide and are the empirical foundation for contemporary commons-based policy — from Barcelona’s digital rights program to the European Union’s Data Governance Act.

A **knowledge commons** is the extension of this idea to non-material resources: software, scientific research, cultural works, data. Knowledge commons are in some ways easier to govern than pasture commons (you can’t overgraze software) and in other ways harder (you can still enclose it, by putting legal fences around it). Most of §4–6 below is about how communities have built legal and institutional fences that keep knowledge commons open.

2. The tragedy of the commons — and why it was wrong

Tier 1

In 1968, the biologist Garrett Hardin published an essay in the journal *Science* called “The Tragedy of the Commons.” His argument was that shared resources are doomed. If ten herders share a pasture, each one benefits individually from adding another cow, but the cost of overgrazing is spread across all ten. So each rationally adds cows, the pasture collapses, and everyone loses. The only way out, Hardin said, is either privatizing the pasture or having a government seize control.

This essay became one of the most cited papers of the twentieth century. It was taught in undergraduate economics, political science, and environmental policy courses for decades. It shaped how a whole generation of policymakers thought about shared resources.

The problem is that Hardin was describing something that isn’t actually a commons. He was describing a free-for-all with no rules. Real commons — the ones that have actually existed in history — have rules, and the rules work. Hardin himself quietly admitted the error in 1994. But by then the original essay had done its work, and the belief that “commons must collapse without centralized control” had become conventional wisdom.

Tier 2

For technologists, this is worth naming because the same assumption keeps coming back in debates about AI safety, internet governance, and data policy. “You can’t have a commons at this scale without central coordination” is the modern form of Hardin’s argument. It’s the unstated premise underneath Karp and Zamiska’s book. It’s also the premise underneath most arguments that AI infrastructure must be concentrated in a few well-resourced firms.

The reason it matters is that every time someone says “commons can’t scale,” they are making an empirical claim that is contradicted by the evidence. Linux runs most of the internet. Wikipedia is bigger than any encyclopedia ever produced. The IETF has governed the internet’s protocol layer through rough consensus for fifty years. These are commons at planetary scale. They work.

Tier 3

For policymakers, the most important fact about the tragedy-of-the-commons frame is that it has been formally refuted by one of the most rigorous bodies of empirical work in twentieth-century social science, and that the refutation carries a Nobel Prize. When you encounter a policy argument that relies on Hardin’s frame — and you will encounter many — you are entitled to ask what the author thinks of Ostrom’s rebuttal. If they have no answer, the argument is forty years out of date.

3. Elinor Ostrom and polycentric governance

Tier 1

Elinor Ostrom (1933–2012) was an American political economist who spent her career asking a simple question: how do real communities actually govern shared resources? She and her collaborators visited farming villages in Switzerland, fishing communities in the Philippines, irrigation cooperatives in Spain, groundwater associations in Southern California, and dozens of other places around the world. They documented how these communities made rules, enforced them, and adjusted them over centuries.

She found eight things that long-lasting commons have in common. Put simply, durable commons have: clear boundaries about who belongs and what's being governed; rules that fit local conditions; ways for people affected by the rules to change them; monitoring by people the community trusts; graduated punishments for rule-breaking; cheap ways to resolve disputes; recognition from higher-level authorities that the community has the right to govern itself; and, for bigger systems, nested layers — smaller commons inside larger ones.

These are not abstract principles. They were derived from watching people actually do this, across hundreds of cases, across centuries.

She called the resulting model **polycentric governance** — many centers of decision-making, coordinating with each other, rather than one central command.

In 2009, she became the first woman to win the Nobel Memorial Prize in Economics, for this work.

Tier 2

For civic technologists, Ostrom's eight principles are not a theoretical framework. They are an engineering specification for durable institutional software. Every one of the eight maps onto concrete architectural choices.

Clear boundaries means your data model has to know which institution owns which data. *Rules fit to local conditions* means configuration has to be local, not global. *Collective choice* means the people affected by a policy need to be able to change it — which means governance features, not just configuration features. *Accountable monitoring* means audit logs that the community can read, not just the vendor. *Nested enterprises* is straight architectural guidance: build systems as nested, federated, with clear interfaces between levels.

The Linux Foundation, Apache Software Foundation, Debian Project, and IETF are Ostromian

in structure whether or not their founders had read her. They evolved toward these principles because the principles are what works.

Tier 3

For policymakers, the policy implications of Ostrom's work are concrete. When governments design infrastructure programs — for digital identity, for health records, for AI — the default assumption has been that scale requires centralization. Ostrom's work shows this is empirically wrong. Polycentric systems are often more resilient, more adaptive, and more legitimate than centralized ones, precisely because they preserve local capacity to adjust.

Estonia's X-Road (see §10 below) is the most important nation-state-scale example: a deliberately decentralized government data architecture, designed in 1999, that has now outperformed every centralized equivalent on every metric that matters — security, adoption, trust, adaptability. It is Ostromian governance operationalized.

When a procurement officer is evaluating whether to put a municipal health system's data inside a single vendor's platform or to build a federated architecture where each institution retains its data, they are making an Ostromian choice, whether they know it or not. The empirical record says the federated model is the durable one.

4. Subsidiarity

Tier 1

Subsidiarity is a political principle that says decisions should be made at the most local level capable of handling them. The town should decide what the town can decide; the state should handle what the town can't; the federal government should handle only what the state can't; and so on. It's the opposite of the principle that everything important should be decided at the top.

The word comes from the Latin *subsidium*, meaning "help" or "support." The idea is that higher levels of authority exist to *help* lower levels do their own work — not to replace them. A national government should help a state government; a state should help a city; a city should help a neighborhood; a neighborhood should help a family. Each level exists to support the level beneath it, not to absorb it.

If this sounds familiar to American readers, it should: federalism is one version of it. But subsidiarity is older than American federalism by a century and a half, and it has been developed more carefully in Catholic social teaching and European political theory than in American constitutional law.

Tier 2

For civic technologists, subsidiarity translates into a clean design rule: **push decisions down, not up**. A building association can make decisions about its own building. It does not need a state-level platform to do that. A school district can make decisions about its own students. It does not need a federal data lake. A city can run its own services. It does not need to route its operations through a vendor's cloud.

The subsidiarity principle is what makes "data sovereignty" a coherent technical goal rather than a political slogan. When we say "the data stays with the institution," we are making a subsidiarity claim: the institution is the right level at which to hold and govern that data, because it is the lowest level competent to do so.

Tier 3

For policymakers, subsidiarity is not only a philosophical principle. It is justiciable European Union constitutional law. Article 5(3) of the Treaty on European Union requires that decisions be made at the EU level only when they cannot be made effectively at the member-state level. The Lisbon Treaty's Protocol 2 established an "Early Warning System" through which national parliaments can formally challenge proposed EU legislation on subsidiarity grounds. This is not a metaphor; it is enforceable law with a case record.

Subsidiarity also has an unusually broad political coalition behind it. Catholic social teaching (Leo XIII 1891, Pius XI 1931, John Paul II) has developed it most thoroughly. American Catholic legal scholars (Robert Vischer, Patrick Brennan) have traced its resonance with federalism. European Christian Democratic parties have built their constitutional programs around it. Communitarian thinkers like Alasdair MacIntyre and conservatives like Robert Nisbet have argued that modern liberal democracies fail precisely because they have eroded the intermediate institutions that subsidiarity presupposes. On the left, Elinor Ostrom's polycentric governance and the distributist economic tradition make overlapping claims.

This matters because the political coalition for distributed institutional sovereignty is wider than the Silicon Valley debate suggests. It includes Catholic social thought, American federalist conservatives, European Christian Democrats, left-communitarian thinkers, commons theorists, and indigenous sovereignty movements. The only group it doesn't include is the one currently writing the AI policy agenda.

5. Ilich and convivial tools

Tier 1

Ivan Illich (1926–2002) was an Austrian-born Catholic priest, philosopher, and social critic who wrote a series of short, disturbing books in the 1970s about how modern institutions capture and disable the very human capacities they claim to serve. His most famous argument, in *Medical Nemesis* (1975), was that past a certain point, the medical system makes people sicker than it makes them well. In *Deschooling Society* (1971), he argued that schools disable the capacity to learn. In *Tools for Conviviality* (1973), he generalized the argument.

Illich's claim was that every tool — every institution, every technology — has two thresholds. Past the first, the tool helps. Past the second, it starts to disable the thing it was supposed to help with. Think of a car: a car extends your ability to travel. But past a certain density of cars, the car disables your ability to walk, because the city has been rebuilt around cars and the sidewalks have been narrowed or removed. "That motor traffic curtains the right to walk," Illich wrote, "not that more people drive Chevies than Fords, constitutes radical monopoly."

The phrase **radical monopoly** matters. Illich wasn't using it to mean a firm being too big. He meant: when a tool forecloses the alternative way of doing the thing entirely. When you can't walk anymore, not because you lack a sidewalk but because the whole architecture of the city no longer admits walking as a practice.

Illich's test for whether a tool is **convivial** — meaning whether it supports autonomous human competence — is whether it can be understood, repaired, and modified by the people who use it, and whether it leaves the practice it augments intact.

Tier 2

For civic technologists, Illich's two-watersheds framework is the single best diagnostic tool for evaluating a platform proposal. Does this system, past a certain point, disable the competence of the institution using it? If the answer is yes, you are looking at a second-watershed tool.

Palantir's Gotham and Foundry are paradigmatic second-watershed tools. An intelligence agency that builds its analytic practice around Gotham has not augmented its analysts; it has restructured the craft of intelligence analysis around the tool. An enterprise that places its operational data inside Foundry's Ontology has not gained analytic capability; it has delegated a core competence to a vendor. When the contract ends, the competence does not come back.

The test for convivial infrastructure is whether the institution using it *retains the capacity to do the work without the tool*. This is not a preference; it is a structural property.

Tier 3

For policymakers, the Illich test suggests a simple procurement question: *if this vendor disappeared tomorrow, could the institution still perform its mission?* If the answer is no, the institution has, by Illich's definition, lost the capacity it thought it was buying.

Most major civic technology procurement does not ask this question. It asks about uptime, price, and feature parity. These are the wrong questions, because they assume the institution will always have the vendor. The Illich question is the right one, because it addresses what happens when the relationship ends — which, eventually, it always does.

6. Free and open-source software, and why licenses matter

Tier 1

Free software — “free” as in freedom, not price — is software whose source code is available to read, modify, and share. The movement started in 1983 when a computer scientist named Richard Stallman announced he was going to build an entirely free operating system (the GNU project) because the software industry was becoming dominated by proprietary, closed-source products.

Stallman's key innovation wasn't the software. It was a legal instrument called the **GNU General Public License**, or **GPL**. The GPL uses copyright law — the same law that normally locks software down — to do the opposite: it requires that anyone who uses or distributes the licensed software must also make their modifications available under the same terms. This is called **copyleft**. It is a kind of legal judo that uses the rules of intellectual property to keep the software permanently open.

The GPL turned out to be extraordinarily powerful. It meant that a community could build software — over decades, across thousands of contributors — with confidence that no single company could later come along and close it off. The Linux operating system, which now runs most of the internet's servers and nearly all of its phones (Android is built on Linux), is licensed under the GPL.

Open source is a related but slightly different movement, which started in 1998 and tends to emphasize the practical benefits of open development (better code, faster bug-fixing, easier collaboration) rather than the ethical argument for user freedom. In practice, the two movements overlap almost entirely, and most software is licensed under one of a handful of well-understood open-source or free-software licenses.

Tier 2

For civic technologists, the license is not a detail. The license is the constitution. It determines what the community can do, what the company sponsoring the software can do,

and what happens when the two come into conflict.

There are two broad families of licenses. **Permissive licenses** (MIT, BSD, Apache) let anyone do almost anything with the code, including embedding it in closed products. **Copyleft licenses** (GPL, LGPL, AGPL) require that modifications stay open. Permissive licenses are friendlier to commercial adoption; copyleft licenses are stronger guarantors of long-term openness.

The **AGPL** (Affero General Public License, 2007) is specifically designed for the cloud era. It closes what used to be called the “ASP loophole” in the regular GPL: under the GPL, if you run modified software on your own servers and only let users access it over the network, you didn’t have to share your modifications. The AGPL closes this loophole. If you run AGPL software as a service, you must share your modifications. This is why AGPL matters now: it is the license that makes a commons durable against hyperscaler enclosure.

When we say Cube Commons uses AGPL and related copyleft licensing, we are making a specific technical and political commitment: the commons we build cannot later be closed off by us or by anyone else. The license prevents it.

Tier 3

For policymakers, the most important thing to understand about open-source licensing is that it works. The free-software movement, over four decades, has built a body of law that courts have enforced, that major companies have learned to respect, and that underlies most of the internet’s actual infrastructure.

What happened between 2018 and 2025 is an important live case. Four commercial companies that had been distributing their software under open-source licenses — MongoDB, Elastic, HashiCorp, and Redis — tried to relicense to more restrictive terms, to prevent Amazon Web Services from reselling their work. Each attempt triggered a community fork: OpenSearch (forked from Elasticsearch in 2021), OpenTofu (forked from Terraform in 2024), and Valkey (forked from Redis in 2024). The Linux Foundation hosted each of the forks, providing neutral governance. Within a year of the Redis fork, 83% of large enterprise Redis users were on or testing Valkey instead.

In 2024 and 2025, Elastic and Redis reversed course and re-adopted open-source licenses — specifically the AGPL. The commons won. The companies that tried to close it lost not only the ideological battle but also, according to an academic study published in November 2024, the financial one: relicensing did not improve their revenue.

This is the live evidence that commons-governed software infrastructure can defeat attempts at enclosure by well-resourced firms. It is the strongest available argument for the Cube Commons model, and it happened in the exact infrastructure layer Cube Commons operates in, in the last five years.

7. Local-first software

Tier 1

Local-first software is an approach to building computer applications that keeps your data on your own devices, not on someone else's servers, while still letting you collaborate with others and sync across devices.

For most of the 2010s, the industry trend went the opposite direction. Your photos went to Google or Apple's cloud. Your documents went to Dropbox or Microsoft 365. Your messages went to WhatsApp's servers. Your health records went to a hospital chain's vendor platform. The benefit was convenience: everything worked everywhere. The cost was that you didn't really own any of it. If the company changed its terms, raised its prices, went out of business, or decided to lock you out, your data went with it.

Local-first software flips this around. Your data lives on your computer, your phone, your organization's own servers. When you want to collaborate with someone or sync across devices, the software does that — but the data doesn't become the cloud provider's property, and the collaboration doesn't require the cloud provider to be available.

Tier 2

The local-first movement was articulated most crisply in a 2019 paper by Martin Kleppmann and colleagues at Ink & Switch, who laid out seven ideals: no spinners (the app is always fast because the data is local); your work is not trapped on one device; the network is optional; seamless collaboration with colleagues; the Long Now (the software keeps working in twenty years, not just while the company behind it exists); security and privacy by default; and ultimate ownership and control.

The technical underpinning is a class of data structures called **CRDTs** (Conflict-free Replicated Data Types), which let multiple people edit the same document on their own devices and then reliably merge the edits without needing a central server to arbitrate. The best-known library is Automerge.

Local-first is not a rejection of collaboration. It is a reorganization of who owns what. Collaboration happens through peer-to-peer synchronization or through thin coordination services — not through a vendor owning a copy of your data.

Tier 3

For policymakers, the local-first architecture answers a question that GDPR and similar data-protection laws have struggled with: how do you give individuals and institutions real

control over their data when all the technical infrastructure assumes the cloud provider has a copy?

The European Union's answer has been to regulate the cloud providers — data protection by design (GDPR Article 25), the Data Governance Act, the Data Act. These are useful but partial. The more structural answer is to build infrastructure that does not assume the cloud provider has a copy in the first place. Local-first software is the architectural expression of the data sovereignty principle.

Estonia's X-Road (see below) is a nation-scale instance of the same principle: each government ministry keeps its own data, and X-Road is the peer-to-peer protocol that lets ministries share what they need to share without anyone having a central copy of everything.

8. Polycentric governance in practice: IETF, W3C, Linux Foundation

Tier 1

Most of the internet's core infrastructure — the protocols that make email work, the protocols that make the web work, the operating system that runs most of its servers — is governed by organizations that do not have CEOs in the normal sense, do not have shareholders, and do not make decisions by executive command.

The **Internet Engineering Task Force (IETF)** defines internet protocols through a process called "rough consensus and running code." That phrase comes from a 1992 plenary speech by David Clark, an MIT computer scientist: "We reject kings, presidents, and voting. We believe in rough consensus and running code." The IETF publishes its decisions as Requests for Comments (RFCs) — the specifications that everyone who builds on the internet implements.

The **World Wide Web Consortium (W3C)** does the same thing for the web. It was founded by Tim Berners-Lee at MIT in 1994, the same year he relinquished all commercial rights to the World Wide Web. CERN, the European physics laboratory where Berners-Lee invented the web, formally released it into the public domain in April 1993. The counterfactual — where the web became a proprietary product — is one of the most consequential near-misses in modern history.

The **Linux Foundation** is the neutral home for many of the world's most important open-source software projects, including Linux itself, Kubernetes, and the forks we mentioned above.

Tier 2

For civic technologists, these organizations are worth studying closely because they solve, in practice, the coordination problem that centralized control promises to solve in theory. They coordinate vast, distributed, competitive actors around shared technical standards — often better than any government or firm could.

The mechanism is Ostromian. There are clear boundaries (what counts as an IETF decision, who can participate). There are local rules (each working group writes its own charter). There is collective choice (working groups review drafts, produce consensus, and occasionally appeal). There is monitoring (implementations have to interoperate or they fail). There are graduated responses to disagreement (document, revise, fork). There are low-cost conflict resolution mechanisms (the IETF appeals process, the W3C's formal objections procedure).

These are not coincidences. They are the principles of durable commons governance, developed independently in the software world because they are what works.

Tier 3

For policymakers, the track record is the argument. The IETF has been governing the internet's protocol layer for forty years. The W3C has been governing the web for thirty. The Linux Foundation hosts projects whose combined economic impact runs into the trillions of dollars. These are among the most successful institutions of the last half-century, by almost any measure.

They are also almost invisible to most policy discourse, because they don't look like the institutions policy-makers are used to. They have no presidents. They have no headquarters (well — they have offices, but that isn't the center of decision-making). They work by rough consensus and running code.

The important question is not whether polycentric governance can scale. It already has. The important question is why policy continues to be designed as though it cannot.

9. Estonia's X-Road

Tier 1

Estonia is a small country on the Baltic Sea. In the 1990s, after regaining independence from the Soviet Union, it had to build a government almost from scratch. In 1999, it made a decision that has turned out to be one of the most important government technology decisions of the last half-century: instead of building one big central government IT system,

it required each ministry to build its own system, and then created a standard way for those systems to talk to each other securely.

The standard is called **X-Road**. It is, in effect, a protocol — a set of rules for how one ministry's system can request specific information from another ministry's system, with strong security and without either ministry having to give up control of its data.

Estonia now consistently ranks among the top countries in the world for digital government. Services like tax filing, health records, voting, and identity verification happen digitally and securely. Estonians spend, on average, a few hours a year interacting with government bureaucracy, where residents of comparably developed countries spend many days. And the architecture is not centralized. Each ministry still owns its own data. X-Road is the peer-to-peer layer that lets them coordinate.

Tier 2

For civic technologists, X-Road is the existence proof. The most common objection to distributed architectures in government IT is that they cannot deliver the integrated user experience citizens expect. Estonia has been delivering that experience, at nation-state scale, for a quarter-century. The ministries own their systems. The protocol does the integration.

The architecture has been adopted and adapted by Finland (Suomi.fi), Iceland, the Faroe Islands, and — significantly — by several developing countries through the MOSIP project. It is open-source and freely available. Any government considering a centralized procurement would, at minimum, be acting in good faith if it first evaluated the X-Road alternative.

Tier 3

For policymakers, the policy lesson of X-Road is that **the procurement choice is architectural, not just commercial**. When a state or city government awards a centralized contract to a vendor, it is not just buying software. It is committing to a governance model — the vendor's — for the term of the contract and usually well beyond it. When it builds or adopts a federated protocol-based architecture, it is committing to a different governance model, one in which the state retains capacity and the participating institutions retain sovereignty.

Estonia's experience suggests that the federated model is not only more sovereign but more effective. That is not what most procurement advisors assume. It is what twenty-five years of working infrastructure demonstrates.

10. The relicensing wave, 2018–2025

Tier 1

Between 2018 and 2025, four well-known software companies — MongoDB, Elastic, HashiCorp, and Redis — each tried the same move. Their products had been released as open source, which had helped them attract users and build communities. But Amazon Web Services (and similar cloud companies) was reselling their software as a service, making money from it without the companies seeing much of that revenue.

So each company, in turn, changed its license to a more restrictive one. The goal was to block the cloud giants from reselling. The effect was that the communities around the software rebelled. Within weeks or months of each relicensing, a community fork appeared — a copy of the software under the old open-source terms, developed by a different group of contributors, hosted by the neutral Linux Foundation. OpenSearch forked from Elastic. OpenTofu forked from HashiCorp's Terraform. Valkey forked from Redis.

The forks grew fast. Within a year of the Redis fork, a large majority of enterprise Redis customers had moved or were moving to Valkey. And then the companies reversed course. Elastic re-adopted an open-source license (AGPL) in 2024. Redis re-adopted AGPL in 2025. HashiCorp was acquired by IBM.

An academic study published in late 2024 found that the relicensing had not even achieved its financial goal: the companies that did it did not see improved revenue.

Tier 2

For civic technologists, this is the most important contemporary case study of commons governance defeating attempted enclosure. It happened in the exact layer where it matters. It happened fast. And the mechanisms that defeated the enclosure are the mechanisms Cube Commons relies on: neutral hosting (Linux Foundation), legal instruments (AGPL, which was designed for exactly this), enterprise procurement pressure, and fast community mobilization.

The case also shows the limits of goodwill. Each of the four companies was, at the time of the attempted enclosure, a beloved participant in the open-source community. Reputation was not enough. What defeated the enclosure was the institutional infrastructure — hosts, licenses, forks — that made the commons *legally* durable against attempted closure.

Tier 3

For policymakers, the implication is direct. When evaluating whether to fund, adopt, or mandate infrastructure, the question is not whether the provider is "good" at this moment. The question is whether the *institutional infrastructure* — licenses, governance, neutral hosts — makes the commons durable against future attempts at closure.

Cube Commons' commitment to AGPL copyleft, to Public Benefit Corporation structure, to protocol-based coordination rather than platform-based coordination, and to institutional data sovereignty is not ornamental. It is the specific set of choices that make the commons we are building resilient against the enclosure pressures that, five years from now, will inevitably come. The relicensing wave is not a historical curiosity; it is the blueprint for what to build against.

11. Distributed institutional sovereignty

Tier 1

Pulling everything together, the phrase **distributed institutional sovereignty** names the architecture we are arguing for. It has three parts.

Distributed means many centers, not one. Real civic institutions — hospitals, housing authorities, schools, libraries, town governments — each hold their own capabilities, each govern their own data, each make their own decisions.

Institutional means that these are actual existing institutions. Not individuals acting alone. Not crypto-secessionist communities exiting the state. Not new charter cities bypassing democratic governance. The existing civic institutions that most of us already belong to and depend on.

Sovereignty means these institutions have real, meaningful control over the infrastructure they rely on. Their data lives with them. Their computation is theirs. The tools they use are tools they can understand, modify, and if necessary walk away from.

This is the architecture that, we think, keeps democratic life working in the age of artificial intelligence. Not by concentrating AI capability in a few state-aligned firms, and not by atomizing it into millions of disconnected individuals, but by strengthening the intermediate institutions that already do the work of a democratic society — and giving them the infrastructure to do that work on their own terms.

Tier 2

The operational shape for civic technologists: the institution's data lives in the institution's own cloud account (or, increasingly, on its own servers). The operator of any shared service operates a separate piece of infrastructure that connects to — but does not contain — the institution's data. The upstream software is open-source and copyleft-licensed.

Coordination happens through open protocols, not proprietary platforms. Governance happens at each layer, with no layer able to capture the others.

The building owns its data. The neighborhood owns its coordination. The city owns its integration. The state owns its regulation. Each level is helped, not absorbed, by the level above. This is subsidiarity operationalized in software.

Tier 3

The political coalition for distributed institutional sovereignty is larger than it looks. Catholic social thought has defended it since 1891. European Christian Democracy has built constitutional programs around it. American federalists have argued for it since 1787. Left communitarians, commons theorists, indigenous sovereignty movements, and digital rights advocates have converged on overlapping versions. Elinor Ostrom won the Nobel Prize for demonstrating it empirically.

What is new is that the infrastructure to make it operational in software now exists — the licenses, the protocols, the architectural patterns, the legal forms. For the first time, a local institution can plausibly run its own digital infrastructure with the same seriousness it runs its physical infrastructure. What is missing is the policy articulation — the clear statement that this is an option, and what it looks like, and why it is the more durable democratic architecture.

That is what “The Other Answer” is for. This primer is its footnotes.

Further reading

For the reader who wants to go deeper, the following are the canonical primary sources behind each concept in this primer.

On commons and polycentric governance: Elinor Ostrom, *Governing the Commons* (Cambridge University Press, 1990); Ostrom’s Nobel lecture, “Beyond Markets and States: Polycentric Governance of Complex Economic Systems” (*American Economic Review*, 2010).

On the failure of the tragedy frame: Garrett Hardin, “The Tragedy of the Commons” (*Science*, 1968), paired with Hardin’s own correction, “The Tragedy of the Unmanaged Commons” (*Trends in Ecology & Evolution*, 1994).

On subsidiarity: Johannes Althusius, *Politica Methodice Digesta* (1603); Pope Pius XI, *Quadragesimo Anno* (1931), §79; Robert Nisbet, *The Quest for Community* (1953); Alasdair MacIntyre, *After Virtue* (1981).

On convivial tools: Ivan Illich, *Tools for Conviviality* (1973); E.F. Schumacher, *Small Is Beautiful* (1973).

On the foundational vision of computing: Norbert Wiener, *The Human Use of Human Beings* (1950); J.C.R. Licklider, "Man-Computer Symbiosis" (1960); Douglas Engelbart, "Augmenting Human Intellect" (1962); Ted Nelson, *Computer Lib / Dream Machines* (1974); Eden Medina, *Cybernetic Revolutionaries* (MIT Press, 2011), on Project Cybersyn.

On free and open-source software: Richard Stallman, "The GNU Manifesto" (*Dr. Dobbs's*, 1985); Eric Raymond, *The Cathedral and the Bazaar* (1999); Lawrence Lessig, *Code and Other Laws of Cyberspace* (1999).

On local-first software: Martin Kleppmann, Adam Wiggins, Peter van Hardenberg, and Mark McGranaghan, "Local-first software: you own your data, in spite of the cloud" (Ink & Switch, 2019).

On Estonia's X-Road: the project documentation at x-road.global; Taavi Kotka's various writings on Estonian digital government.

On the relicensing wave: Brockmeier et al., "An Analysis of Open-Source License Changes" (CHAOSS/OpenForum Academy, November 2024, arXiv:2411.04739).

On platform concentration and its critics: Shoshana Zuboff, *The Age of Surveillance Capitalism* (2019); Cory Doctorow, *The Internet Con: How to Seize the Means of Computation* (2023); Ben Tarnoff, *Internet for the People* (2022); Meredith Whittaker, David Gray Widder, and Sarah West, "Open (For Business)" (2023, forthcoming in *Nature* 2024).

On the book we are responding to: Alexander C. Karp and Nicholas W. Zamiska, *The Technological Republic: Hard Power, Soft Belief, and the Future of the West* (Crown Currency, 2025).

Cube Commons, Inc. is a Massachusetts Public Benefit Corporation building local-first, open-source multi-agent coordination infrastructure for civic institutions. This addendum accompanies "The Other Answer" (CC-TR-2026-005) and is deposited on Zenodo under CC BY 4.0.