



The blueprint

North Canterbury · © John Stroh

Version 0.2 · revised 24 August 2026 · [this version as a PDF](#)

Drafted with AI assistance, then checked and revised by the author. The judgements and the errors are the author's.

Status of these claims

What this publication does not claim, and what is outstanding against it in the register.

A question this rests on is parked: What counts as complete state for a cluster?

We do not claim a defined completeness boundary for cluster state; we claim the criteria in DEL-06 and nothing beyond them.

A question this rests on is parked: What secession premium will members bear?

We publish no pricing and no affordability claim, and we do not assert that members value the secession guarantee.

A question this rests on is parked: What do we do with an application that refuses to run inside the boundary?

We do not claim that a member's existing applications can be made to run inside the boundary. We claim only that the boundary reveals which ones cannot.

A question this rests on is open: Does secession capability actually convert a buyer?

What a group of small organisations builds together, and what it has to do.

Series part C · Version 0.2 · August 2026

Before you read this

You have read part A and recognised your own situation in it. Four questions follow immediately, and this part answers them.

What exactly are we building? Not a buying club and not a better vendor. Something more specific, and the specifics matter because most of them are what stop it turning into the thing you were escaping.

What does it actually have to do? There is a short list. Twelve capabilities, each of which can be checked by watching it work. Miss one and a particular question about your own organisation becomes unanswerable.

What do we have to decide ourselves? More than you would like. Some of it we cannot decide for you and some of it nobody can decide in advance.

And how does this fail? It fails in four ways that are well understood, and knowing them in advance is most of the defence.

Nothing here asks you to buy anything. If you build this and never make contact with anyone who wrote it, that is the intended outcome.

What you are building

A **co-operative** that owns the infrastructure its members depend on.

Not a group discount. A group discount leaves the same relationship in place with a better price, and the price goes back up when the group weakens. What changes the situation is ownership: the terms cannot be altered against you by someone whose interests differ from yours, because nobody is in that position.

Four properties do the work, and each of them is a decision that has to be made once and written where it cannot easily be undone.

Members hold control, not investors. Under co-operative company law in this country, transacting members must hold at least sixty per cent of voting rights. That floor is the mechanism by which everything else stays true. Investors may hold shares; they may not hold

control.


Members own their own machines outright. No licence key, no remote disable, no security taken over the equipment. The moment the co-operative can switch off a member's hardware, it has recreated the dependency it exists to prevent. Where a member is financed, the security is their share, never the box.

The revenue may not come from member data. No advertising, no brokerage, no arrangement where a third party pays for access to what members hold. In a privacy policy this is worth nothing. In the constitution, amendable only by the members it protects, it holds.

A region can leave and keep working. Any cluster can separate from the rest with its data intact and continue operating. Not asserted — exercised annually, with a member watching, and the result published whether it passes or fails.

That last one is the guarantee that makes the rest credible. Any organisation can promise the first three. Only one that has demonstrated the fourth has shown that the promises are enforceable by the people they were made to.

FIG-04 · QUERIES FEDERATE · RECORDS DO NOT



CON-07 · drilled annually, member observing, result published whether it passes or fails.

FIG-04 Federation is an enrichment layer over complete clusters, never a dependency of them. A cluster can be severed and keep operating.

What it has to do

Twelve capabilities. They are specified as **properties, not methods** — what the thing must do, never how it must be built — so that more than one design can satisfy them and no supplier is written into the standard.

Each is verified by watching it happen, not by reading a description of it.

	CAPABILITY	VERIFIED BY
1	Inference on hardware members control	Continued operation for routine work with the external network unreachable
2	Egress control with signed attestation	Twelve months of boundary-derived attestations a member can verify independently
3	A substitutable inference endpoint	A third-party application, not written for this environment, running against a local endpoint with no egress exception
4	Supply-chain declaration	A register naming every component that performs inference, with operator and jurisdiction, and a notification record
5	Per-record authority, default deny	Enforcement failing closed when the grant store is unreachable; a member listing every grant they hold
6	Provenance on derived artefacts	A query returning, for any record, everything derived from it — run, not described
7	Collective authority	An administrator attempting to release a jointly held record, and failing
8	No composition of a person	The absence of any operator-side path by which records about one person can be assembled across sources
9	Durable signed records	A record written before a signature-algorithm change still verifying after it
10	Complete-state export	A timed export and restoration, performed by the people who would actually perform it
11	A saleable artefact cannot be inverted to a record	An adversarial attempt, by a party who did not build it, to recover an identifiable member record from the artefact — performed and failed, annually, result published either way

	CAPABILITY	VERIFIED BY
12	An agent cannot widen its own authority	An agent instructed explicitly to widen its own grant, attempting it, and failing — with the attempt in the audit trail

Capability 11 — three conditions, all of which must hold before an artefact is saleable

1. **Authored, not held.** The artefact is produced by the co-operative. No member can point to it and say *that is mine*.
2. **Inversion resistant.** The adversarial attempt fails. Not asserted — attempted, by someone with an interest in succeeding, on the same footing as the CON-07 exit drill. A member observes. The result is published whether it passes or fails.
3. **Provenance complete and granted.** Capability 6 as written runs forward — for any record, everything derived from it. This condition needs it run backward, from an artefact to every member whose contribution entered it, and that inverse is not yet specified. When it is, the query returns every such member, and each of those members holds a live, revocable grant covering that class of derivation. Default deny. A revoked grant means the artefact is withdrawn from sale at the next release, not argued about.

An artefact failing any of the three is not sold. Not sold pending review — not sold.

Capability 12 — four properties


1. **The grant store is not agent-writable.** Capability 5 already provides per-record authority with default deny and fail-closed behaviour when the grant store is unreachable. The addition is that no agent may mint, widen, extend or self-issue a grant. Grants originate from members. Full stop.
2. **Every agent action names the grant it acted under.** Attribution is read, not assigned afterwards. An action with no grant id is not an unattributed action — it is a refused action.
3. **Capability changes are versioned and reversible.** A model, index or tool change that degrades behaviour must be revertible to a known state. Verified the same way as capability 10: a timed rollback, performed by the people who would actually perform it.
4. **Domain expansion requires a vote.** The agentic layer is where this will be tested, because the pressure to widen a grant always arrives attached to a task someone wants done today.

Rider. The minimum quorum for a grant widening is unspecified, and whether it differs from the quorum for a constitutional change is unsettled. The answer determines how fast the system can be useful. Recorded here rather than assumed.

Numbers 1 and 3 are the pair that matters most, and 3 is the one usually missed. Most AI reaches an organisation inside software it already bought. A system that can only serve applications written for it addresses a small and shrinking part of your actual exposure.

Number 8 is the one that will be hardest to hold. Assembling a working picture of a person from their mail, calendar, documents and records is not a misuse of an agent — it is what an agent is for. Refusing it is a deliberate limitation, and it has to be architectural rather than a policy, because a system that merely retains the ability to do it will eventually do it.

FIG-11 THE BOUNDARY YOU MANAGE IS A BOUNDARY THE RECORD ALREADY HAS



A handoff that carries no evidence is pure cost. Draw the boundary where the record already draws one.

FIG-11 Functional partitions require operations to tell governance what happened, and agent event volume outruns the telling. Authority partitions put the boundary where the record already has one.

What you have to decide

Eight things. We have views on some and no standing on any.

Who anchors it. Public bodies have budget and procurement pain but want assurance before they will pilot, which is real cost before any revenue. Community organisations move faster and pay less. Our view is that public-sector members are the anchor, because their requirements are the ones that force the standard to be real. It is a defensible decision either way and it shapes everything downstream.

How big a cluster is, and where. Small enough that members recognise each other; large enough that the hardware is not idle. Regional, or organised around a shared purpose. The trade-off is utilisation against trust, and only you know which is scarce where you are.

● **This section is blocked by CON-04 and is not settled.** The non-extraction clause reads: *“No part of the co-operative’s revenue may derive from the sale, brokerage, or use for third-party benefit of member data **or artefacts derived from it.**”* Derived artefacts are named in the prohibition. Everything below about customers buying artefacts, and about capability 11 making an artefact “saleable”, therefore describes a revenue model this co-operative’s own constitution forbids as it currently stands. **Either CON-04 is amended by the members, or these claims come out.** That is a constitutional decision and it belongs to the members, not to this document. It is stated here rather than left for a reader to find, because the paragraph two screens up paraphrases CON-04 as *“no arrangement where a third party pays for access to what members hold”* — which drops the five words that decide it.

Who is a member, and who is a customer. They are not the same and the blueprint had only one word for both. **A member** transacts, governs, contributes, holds records inside a cluster, and bears the exit obligation — membership stays narrow and requirement-forcing. **A customer** buys derived artefacts: no governance rights, no records inside a cluster, no vote. Revenue widens without the electorate widening, and the sixty per cent transacting-member floor stays comfortable rather than becoming something to manage.

Two hazards come with that split, and both need a number rather than a policy. **Customer capture** — a customer large enough becomes a governor in fact whatever the constitution says, because the co-operative cannot afford to displease it. The defence is a published concentration figure: revenue by customer, annually, with a cap the members set, in the annual report where a member can check it. **Query-as-a-service** — the obvious commercial temptation is to sell customers a query interface rather than an artefact, and a query answered against held records is brokerage no matter how the invoice reads. Customers buy artefacts that have passed capability 11. They do not buy questions.

How contributed work converts to entitlement. Field service performed by accredited members rather than bought is the difference between a cost model that holds and one that inverts, so it needs specifying rather than assuming. **Accredited against what:** competence per task class — hardware replacement, index rebuild, drill observation, grant custody, corpus curation, first-line service response — with a register naming who is accredited for what, visible to the cluster. **Certified by whom:** if the executive certifies, it controls the currency and the co-operative has recreated the dependency it exists to prevent; peer certification within a cluster, audited across clusters, is the shape most likely to hold. **Converted at what rate:** a rate the board can change is a currency the board controls, so the rate is set by member vote, published, and changed only prospectively — entitlement already earned is never retrospectively revalued. **Lapsing, or not:** accreditation lapses, because competence decays and the register must reflect that; entitlement earned does not. Conflating the two would let the co-operative claw back membership from someone who stopped turning up, which is a different and much worse institution.

The number that tells you whether this is working is the **service coverage ratio**: the proportion of field service performed by accredited members rather than bought, measured per cluster and published annually. It is the early warning for the support-cost failure below — a cluster whose coverage ratio is falling is a cluster whose cost model is inverting, and it shows up here a year before it shows up in the accounts.

Rider. Whether the entitlement mechanism is lawful in this form has not been tested. One legal opinion decides whether it exists at all, and it has not been obtained. What is set out above is the shape the mechanism would take; it is not advice that the shape is available.

What a share entitles you to. Our recommendation is that shares are measured in **capability rather than dollars** — a share buys a defined tranche of capacity, and the dollar price moves as hardware costs fall. That way the member's return arrives as more capability for the same holding, membership never becomes an investment, and later members are not priced out. It is unusual and it will need a lawyer.

Whether contribution earns membership. Ours does: accredited hours, domain expertise and shared procedures earn entitlement on the same footing as cash. Without it you exclude every organisation rich in knowledge and short of money, which is most of the ones this is for. With it, contributing knowledge stops being extraction and becomes investment in something the contributor owns a share of.

What you will not do. Write it down early. The list of things the co-operative is constitutionally barred from is more valuable than the list of things it offers, because the second changes and the first is what members joined for.

And what it costs. We cannot tell you, and part F sets out the framework rather than a figure. The number that matters is **cost per site to install and sustain for five years**. A minister and a board of trustees need the same figure and everything else ladders to it.

How to start

Not nationwide, and not by forming the co-operative first.

One cluster, where trust already exists. A group that already works together — a professional network, a set of community organisations, a district. Prove the whole loop end to end: hardware installed, models running locally, the service desk answering, records signed, the exit drill passed.

Then form the co-operative around the people who have used it. Co-operatives formed by organisations that have already transacted together hold. Ones formed on a prospectus do not.

Then the second cluster, deliberately different. Different hardware generation, different base model. This will feel like an unnecessary expense and it is the single most important structural decision after ownership: a federation running one configuration fails everywhere at once, and the pressure to standardise will never stop arriving with a good argument attached.

How this fails

Four ways, all documented, none mysterious.

Support costs more than you modelled. The economics rest on routine questions being answered locally and people being needed only by exception. If that ratio is worse than assumed, or arrives two years later than assumed, the cost model inverts. Design the escalation rate as a measured, published number and make the business case survive three times your target. If it only works at five per cent escalation you are betting the organisation on a capability forecast.

Field service breaks you. Hardware in a hundred buildings is a field service business, and field service destroys organisations that did not plan to be in it. Ruthless standardisation is the usual answer and it conflicts directly with the heterogeneity requirement above. The resolution is standardising within a cluster and differing between them.

Member capital is slow. Co-operatives grow at the speed of member trust, which is the right speed and not a fast one. Budget for it. Do not promise a national footprint in year two.

Founder dependency. If the whole thing depends on one person's judgement, it is not sovereign either — it has a single point of failure with a name. Write down how hard questions were decided and why, from the first one, and make adding to that record an obligation rather than a courtesy.

How you would know this is wrong

Three things would falsify the proposition rather than merely complicate it.

1. **Members do not value the exit guarantee enough to pay for it.** Holding complete state at every cluster costs more than centralising. If most members would trade the guarantee for a lower price, the design is wrong and should be rewritten rather than better argued.
2. **The routine-locally ratio does not hold.** If the share of work a local model can handle is materially lower than assumed, the service economics fail regardless of governance.
3. **Existing software refuses to run inside the boundary at scale.** If nearly every application a member depends on requires offshore inference and cannot be pointed elsewhere,

the discovery capability is still useful but the substitution one is not.

All three are measurable. If you build this, measure them early and publish what you find, including if it goes against you.

Published under CC BY 4.0. Copy it, adapt it, argue with it. Nothing here names a supplier, and nothing in it requires any particular organisation to exist.

Cite this version

The blueprint, version 0.2 (August 2026).
`agenticgovernance.digital/downloads/the-blueprint-v0-2.pdf`

Licensed CC BY 4.0. Reuse is free, including commercially, provided the author is credited and the reuse does not suggest that the author endorses it. This link is frozen at version 0.2; [/the-blueprint](#) always shows the current one.

PREVIOUS

B · Six questions worth asking

Series contents

NEXT

D · What it does for you on a Tuesday

Alongside: [questions and answers](#) · [glossary](#) · [sources and provenance](#) · [slides](#)