

Venture Investment Memorandum

Evidence Infrastructure for Agricultural Climate Reliance

Australia-first evidence infrastructure for agricultural climate reliance

Investment thesis. Commercialize the working Evidence Event Inbox as a narrow Australian market-entry wedge: map real agricultural operational evidence into an inspectable and portable evidence contract; carry it through a bounded external review or reliance artifact; and expand only when contracted revenue, technical reuse, and documented external reliance justify the next stage.

Prepared for	Prospective Seed Investors and Strategic Venture Partners
Investment proposition	Finance a milestone-gated seed program to convert a working evidence bridge into a repeatable, high-margin infrastructure company
Market focus	Australian livestock technology, methane evidence, agricultural Scope 3 and inset-ting, assurance, and institutional reliance
Commercial posture	Developer-led, evidence-first, no replacement of operational systems of record
Working repository basis	Commit f71d04d16c747374c215734e8ce191824dd63592 as the Event Inbox com-mercial spine; 1b9ff053b51b0dda2ed14a19e0f491e69629e9e4 as the campaign and self-service layer; exact external baseline to be verified and tagged before launch
Primary currency	U.S. dollars; account-specific Australian-dollar models retained as secondary inter-nal diligence inputs
Program horizon	30-day proof; 90-day commercial test; 12-month revenue-gated productization; 24-month publication and institutional adoption sequence
Status	Confidential fundraising memorandum; product, customer, pricing, and revenue assumptions remain subject to diligence
Date	August 2026

This memorandum consolidates the supplied product outline and AgEvidence Australia strategy materials into a venture-investor narrative focused on category creation, product leverage, market entry, milestones, and diligence. All named-account, funding, regulatory, repository, scientific, pricing, and institutional claims require dated verification before external use. The unrelated consumer-packaged-goods summary is excluded from the core product thesis.

Investment Thesis at a Glance

Venture thesis

AgEvidence is building the evidence and reliance infrastructure layer beneath agricultural climate markets. The company does not replace farm, livestock, telemetry, laboratory, processor, remote-sensing, or supply-chain systems. It converts their existing operational records into inspectable, portable, policy-aware evidence artifacts that can be used by buyers, verifiers, registries, research sponsors, lenders, and reporting entities.

The seed-stage opportunity is to prove a narrow but expandable wedge: paid mapping of livestock operational evidence into a reusable evidence contract, followed by production hardening, bounded external review, and recurring managed evidence operations. Venture-scale potential depends on converting customer-specific work into shared schemas, adapters, conformance assets, artifact profiles, and operating infrastructure that lower marginal delivery cost over time.

Investor question	Current thesis	Milestone evidence required
Why now?	Climate claims increasingly depend on upstream operational evidence that existing systems do not package for institutional reliance.	Named external obligations, reviewer demand, and paid evidence-mapping engagements.
Why this wedge?	The working Event Inbox and source-manifest paths preserve customer systems while exposing provenance and evidence gaps.	One end-to-end real mapping, one external-review artifact, and a second-customer reuse test.
What is the product?	A thin-waist evidence contract, trust-state model, graph architecture, policy adapters, artifact builder, and managed evidence plane.	Production hardening, reproducibility, security controls, reusable adapters, and service-level performance.
What is the business model?	Bounded implementation and reliance engagements that expand into recurring managed evidence operations.	Contracted revenue, gross-margin progression, lower support effort, and annual recurring revenue.
What is the moat?	Accumulated evidence grammar, source mappings, reusable adapters, conformance fixtures, reviewer expectations, scientific publication, and reliance history.	External implementations, independent citations, repeated institutional use, and declining bespoke code per customer.
What could invalidate it?	Customers may prefer generic carbon accounting, bespoke consulting, or existing assurance workflows.	Explicit stop conditions at Day 30, Day 90, and Month 12.

Non-negotiable guardrail

The repository materials describe a working scaffold for signed event intake, evidence projection, receipt requests, replay, webhooks, dead-letter handling, source manifests, sandbox quotes, artifact orders, and local verification. They do not establish production security, scientific efficacy, method eligibility, regulatory or registry acceptance, customer adoption, institutional reliance, payment collection, or recognized revenue. These are diligence items and milestone gates, not completed proof points.

Investor Reading Guide

Reader	Priority sections
Venture investors	Sections 1–7, 15–21, and 24–27: investment thesis, market wedge, business model, milestones, economics, risks, and diligence.
Product and engineering	Sections 8–14 and Appendices A–C: architecture, graph model, trust states, modules, capability register, data contract, and acceptance criteria.
Commercial diligence	Sections 5–7 and 15–20: target market, customer qualification, packaging, pricing, conversion, revenue stages, and paid pilot design.
Science, assurance, legal, and data	Sections 9–14, 22, 24–26: state separation, rights, method adapters, review, publication independence, claims controls, and data governance.
Strategic investors and ecosystem partners	Sections 19, 22, and Appendix D: open evidence commons, conformance assets, publication sequence, and institutional distribution.

Contents

Investment Thesis at a Glance	i
Investor Reading Guide	ii
1 Investment Memorandum Purpose	1
2 Investment Summary	1
2.1 Proposed financing and milestone use	1
2.2 Product thesis	2
2.3 Strategic outcome	2
2.4 What success looks like	2
3 Venture Thesis: Why This Wedge, Why Now	2
3.1 The unresolved proof is commercial and institutional	2
3.2 Three strategic options	3
3.3 Why Australia first	3
3.4 Why livestock first	3
4 Product Vision and Positioning	3
4.1 Vision statement	4
4.2 Positioning statement	4
4.3 The thin-waist model	4
4.4 Product principles	4
5 Customer Problem	5
5.1 Current alternatives and their limits	5
6 Target Market and Initial Account Sequence	6
6.1 Initial market boundary	6
6.2 Primary customer groups	6
6.2.1 Event-first technology providers	6
6.2.2 Source-record-first companies	6
6.2.3 Measurement and processor platforms	7
6.2.4 Corporate buyers and program operators	7
6.3 Wave 1 target sequence	7
6.4 Target-account qualification	7
7 Customer Jobs to Be Done	8
7.1 Job 1: Connect evidence without replacing the source platform	8
7.2 Job 2: Establish provenance and integrity	8
7.3 Job 3: Identify evidence gaps before formal review	8

7.4	Job 4: Evaluate the same evidence under multiple policies	8
7.5	Job 5: Assemble bounded reliance artifacts	8
7.6	Job 6: Operate recurring evidence flows	8
8	Product Architecture Overview	8
8.1	Architecture thesis	8
8.2	Ingestion path A: signed operational events	9
8.3	Ingestion path B: controlled source records	9
8.4	Evidence processing lifecycle	9
9	Canonical Graph Model	10
9.1	Graph invariants	10
9.2	Minimum entity relationships for the first livestock wedge	11
10	Trust-State Model	11
10.1	Required state-transition controls	12
11	Core Product Modules	12
11.1	Evidence Event Inbox	12
11.2	Source Manifest Manager	12
11.3	Evidence Projection Engine	13
11.4	Evidence Gap Analyzer	13
11.5	Human Review Ledger	13
11.6	Policy and Method Adapter Layer	14
11.7	Claim and Rights Engine	14
11.8	Reliance Artifact Builder	14
11.9	Reliance Registry	14
11.10	Managed Evidence Operations	15
12	Capability Register: Implemented, Partial, and Absent	15
12.1	Release truth requirements	16
13	Data Contract and Reliance Artifact Anatomy	16
13.1	Minimum evidence envelope	16
13.2	Reliance artifact structure	17
13.3	Artifact acceptance criteria	17
14	Security, Privacy, Rights, and Operations Requirements	18
14.1	Security requirements	18
14.2	Privacy and data stewardship	18
14.3	Operational requirements	18
15	Product Packaging	19
15.1	Packaging refinement	19
15.2	Standard deliverable definitions	20
16	Pricing Methodology and Commercial Controls	20
16.1	Value-based fee ceiling	20
16.2	Four independent commercial caps	21
16.3	Commercial state discipline	21
16.4	Quote approval packet	21
17	Go-to-Market Model	21
17.1	Zero-sales-team activation	22
17.2	Conversion sequence	22
17.3	Outreach proposition	22
17.4	Private-before-public sequence	22
17.5	Sales boundary	22

18 Open-Core and Commercial Boundary	23
18.1 Open evidence commons	23
18.2 Paid commercial layer	23
18.3 Boundary test	24
18.4 Licensing and governance decisions	24
19 Customer Qualification and Account Diligence	24
19.1 Five-dimensional diligence	24
19.2 Account evidence packet	25
19.3 Scoring model	25
19.4 Disqualification conditions	25
20 Ninety-Day Commercial Program	26
20.1 Program objective	26
20.2 Day 0–30: proof phase	26
20.3 Day 31–60: first implementation and independent reconstruction	26
20.4 Day 61–90: external artifact and reuse proof	26
20.5 Day-30 gate	27
20.6 Day-90 gate	27
21 Twelve-Month Productization Roadmap	27
21.1 Month 4–6 product priorities	28
21.2 Month 7–12 product priorities	28
21.3 Explicitly deferred domains	28
22 Scientific Publication and Institutional Strategy	29
22.1 Objective	29
22.2 Publication themes	29
22.3 Strategic effects	29
22.4 Publication and product synchronization	29
22.5 Scientific independence	30
22.6 Twenty-four-month sequence	30
23 Success Metrics and Management Scorecard	30
23.1 Commercial metrics	30
23.2 Developer and activation metrics	31
23.3 Product and operations metrics	31
23.4 Reliance metrics	31
23.5 Strategic and publication metrics	32
23.6 Thirty-day management targets	32
23.7 Twelve-month directional outcomes	32
24 Governance, Ownership, and Decision Rights	33
24.1 Core roles	33
24.2 Proposed RACI	33
24.3 Decision forums	35
24.3.1 Weekly commercialization review	35
24.3.2 Architecture review	35
24.3.3 Scientific and claims review	35
24.3.4 Data, rights, and security review	35
24.3.5 Commercial approval	35
24.4 Issue taxonomy	35
25 Risks and Controls	35
25.1 Overclaiming readiness	36
25.2 Validity laundering	36
25.3 Building before demand	36
25.4 Consultancy trap	36

25.5 Open-source value leakage	36
25.6 Institutional or scientific capture	36
25.7 Customer data and rights risk	37
25.8 Target-account assumption risk	37
25.9 Policy and adapter drift	37
25.10 External reliance concentration	37
25.11 Operational and security failure	37
26 Stop Conditions	37
27 Investment Case and Financing Milestones	38
27.1 What seed capital must prove	38
27.2 Use-of-funds logic	39
27.3 Financing conditions and investor diligence	39
28 Source Reconciliation Required Before External Finalization	39
28.1 Repository baselines	39
28.2 Campaign versions	40
28.3 Pricing frameworks	40
28.4 Currency	40
28.5 Generic CPG summary	40
28.6 Named-account assumptions	40
28.7 External legal and regulatory facts	40
A Ninety-Day Acceptance Criteria	41
A.1 Release acceptance	41
A.2 First customer sprint acceptance	41
A.3 Day-90 continuation acceptance	41
B Product Requirements Traceability Matrix	42
C Wave 1 Activation Plays	44
C.1 DIT AgTech	44
C.2 MEQ Solutions	44
C.3 Agscent	44
C.4 Cibo Labs	44
C.5 Rumin8	44
C.6 Sea Forest	45
D Source Basis and Claims Register	45
D.1 Source register	45
D.2 Claims classification	45
D.3 Core prohibited claims	46
E Glossary	46
F Final Recommendation	47

1 Investment Memorandum Purpose

This memorandum presents **AgEvidence Australia** as a standalone venture-scale evidence-infrastructure company for livestock methane, agricultural Scope 3 and inseting claims, verification readiness, and bounded institutional reliance. It is designed to support seed-stage investment diligence rather than an internal operating approval.

The initial commercialization strategy is intentionally narrow. AgEvidence should use the implemented Evidence Event Inbox and controlled source-manifest path as an Australian market-entry wedge. The first proof is not a large deployment, a complete carbon-accounting platform, or a standards victory. The first proof is that one or two livestock-technology companies will pay AgEvidence to:

1. identify a material operational evidence flow already generated by the customer's system;
2. map that flow into AgEvidence's evidence contract without replacing the source platform;
3. expose provenance, identity, calibration, authority, method, review, or allocation gaps;
4. assemble a bounded artifact for a named external reviewer, verifier, buyer, sponsor, registry, lender, or reporting entity; and
5. produce reusable schemas, handlers, fixtures, profiles, and operating procedures that reduce the cost of the next engagement.

The financing thesis therefore funds a **milestone-gated commercialization experiment around an existing technical spine**. It does not authorize a greenfield rewrite. It does not assume that every current module is production-grade. It does not infer customer demand from funding announcements, repository activity, sandbox quotes, or scientific interest. It requires demand to be proven through contracted work, external review, and reuse.

Boundary of the investment case

The venture case is the bridge-first product test, not a broad platform replacement, a general-purpose agricultural carbon ledger, a customer-specific country fork, a large sales organization, or expansion into soil carbon, autonomous machinery, fermentation, hydrogen, or other domain families before the livestock evidence wedge proves recurring value.

Source basis: S1 Event Inbox C-Suite Decision; S4 Canonical Zero Sales Team Campaign; S8 Comprehensive Product Proposal Outline..

2 Investment Summary

2.1 Proposed financing and milestone use

The proposed seed financing should fund an Australia-first commercialization program with four linked milestone horizons:

- ✓ **30-day technical and commercial proof.** Freeze and reproduce the working product surface, activate a small private cohort, receive at least two real evidence mappings, issue one to three bounded proposals, and close at least one paid engagement or stop.
- ✓ **90-day bridge-first market test.** Harden one customer source integration, complete one method or institutional profile, assemble one external-review artifact, obtain structured external feedback, and demonstrate reuse with a second customer.
- ✓ **12-month revenue-gated productization.** Launch managed evidence operations only after contracted demand; establish production security and observability; add reusable livestock adapters and artifact profiles; convert at least one project into recurring revenue.
- ✓ **24-month publication and institutional adoption sequence.** Publish the open reference architecture, conformance fixtures, reproducibility assets, and peer-reviewed work while preserving independent scientific governance and the paid commercial layer.

2.2 Product thesis

AgEvidence should become the evidence layer beneath existing agricultural systems, not another system that requires customers to abandon farm-management, telemetry, sensor, remote-sensing, biological-input, trial, processor, or supply-chain platforms.


The product receives a narrow set of operational events or controlled source manifests; preserves source-system authority and provenance; projects records into append-only evidence structures; identifies missing evidence; records immutable human review; evaluates the same evidence under versioned country, method, or institutional profiles; assembles portable reliance artifacts; and records any external decision separately.

The commercial unit is not a seat, login, API token, receipt count, or byte. It is a **bounded evidence-and-reliance engagement** tied to a real external decision and a reusable product increment.

2.3 Strategic outcome

The proposal combines two reinforcing wedges:

1. **Near-term commercial wedge:** paid evidence architecture, implementation, verification readiness, reliance artifacts, and managed evidence operations.
2. **Long-term institutional wedge:** open schemas, reproducible reference implementations, developer adoption, scientific publication, and institutional convergence around compatible evidence contracts.



This flywheel is only credible if the open layer makes adoption easier and the paid layer captures customer-specific implementation, operation, risk, and reliance work.

2.4 What success looks like

At Day 90, the program should be able to show:

1. one real Australian or adjacent livestock company has mapped a material source flow into the evidence contract;
2. one paid scope is tied to an exact repository release, named deliverables, acceptance tests, payment terms, and an external reliance question;
3. one external reviewer can reconstruct the evidence and limitations without undocumented analyst knowledge;
4. a second customer reuses a material share of the first customer's schemas, handlers, projections, fixtures, or artifact profile;
5. technical, scientific, method, verification, reliance, allocation, payment, and revenue states remain separate; and
6. the company and investors have evidence to continue, narrow, or stop without relying on optimistic pipeline language.

Source basis: S1, S2, S4, S6, S8..

3 Venture Thesis: Why This Wedge, Why Now

3.1 The unresolved proof is commercial and institutional

The source materials describe a repository state in which the architectural question has a working answer: an upstream platform can submit a controlled event or source reference into an append-only evidence path without migrating its operating system. The unresolved question is whether that bridge reduces enough verification, diligence, reconstruction, and institutional-friction cost that a customer will pay and an external institution will use the

resulting artifact.

The correct venture-building response is therefore not to build the complete imagined platform before engaging the market. It is to convert the existing bridge into the smallest defensible commercial product, measure the resulting evidence gaps, and let contracted obligations govern hardening.

3.2 Three strategic options

Option	Advantages	Failure mode	Venture assessment
Broad platform build first	Maximum conceptual scope; unified vision; more features available at launch.	High burn, slow learning, speculative domains, customer replacement risk, and no proof of willingness to pay.	Not the seed strategy
Consultancy-only evidence projects	Fast custom revenue; flexible delivery.	Bespoke services trap, weak reuse, no protocol distribution, and limited recurring economics.	Reject as primary model
Bridge-first product test	Builds on working scaffold, preserves source systems, makes gaps visible, ties hardening to contracts, and creates reusable evidence assets.	Requires strict scope and claims discipline; first project may reveal hard production gaps.	Preferred seed wedge

3.3 Why Australia first

The source corpus frames Australia as a concentrated test environment for livestock methane, agricultural Scope 3 and inseting, processor evidence, external assurance, and multi-party claim allocation. The strategic value is not merely geographic. Australia can serve as the first production country adapter for a globally stable evidence core:

- source-system and evidence-envelope semantics remain globally stable;
- Australian applicability, methodology, baseline, uncertainty, assurance, data, and claim rules are expressed as versioned profiles;
- institutional decisions remain local and bounded; and
- the same evidence graph can later be evaluated under other country, buyer, registry, or reporting profiles without rewriting its history.

3.4 Why livestock first

Livestock technology offers the best initial fit because the source documents identify current event and source-record pathways for intervention, dosing, measurement, calibration, model outputs, processor records, trial data, and producer payment. Other domains may be attractive, but they require additional vocabularies and conformance suites. Allowing those domains to shape the first 90 days would dilute the test.

Non-negotiable guardrail

Australia-first does not mean country-specific cryptography or an Australian fork of the global evidence core. Livestock-first does not mean AgEvidence endorses a methane technology or claims efficacy. The product remains technology-neutral and records evidence, method compatibility, review, and external reliance as separate determinations.

Source basis: S1, S4, S5, S6, S8..

4 Product Vision and Positioning

4.1 Vision statement

Capture agricultural evidence once, preserve its provenance once, and make it portable across methods, institutions, buyers, and reporting contexts without rewriting the underlying operational system.

4.2 Positioning statement


For agricultural technology companies facing buyer, verifier, regulatory, registry, research, or financing scrutiny, **AgEvidence** is an evidence and reliance infrastructure layer that transforms existing operational records into inspectable, portable, policy-aware evidence artifacts.

Unlike generic carbon-accounting software, consultancy-only audit preparation, or a replacement farm platform, AgEvidence:

- preserves the authority of upstream source systems;
- commits to events and source manifests without indiscriminately centralizing sensitive files;
- separates technical integrity from scientific and institutional determinations;
- supports repeatable evaluation under versioned methods and policies;
- records human review and external reliance as explicit events; and
- creates a recurring managed-operations path after the first implementation.

4.3 The thin-waist model

AgEvidence should function like the narrow middle of a protocol stack. Many systems can operate above and below a small common contract without sharing the same complete software stack. The evidence contract becomes the interoperability waist:



The product is strategically stronger when many source systems can emit compatible evidence, many methods can interpret it, and many institutions can receive bounded artifacts, while no participant must adopt a monolithic AgEvidence application.

4.4 Product principles

1. **Preserve upstream authority.** AgEvidence records and projects evidence; it does not silently become the customer's operational system of record.
2. **Append, do not rewrite.** New evidence, review, policy, allocation, payment, or reliance events extend history.
3. **Separate state families.** Integrity is not science; science is not method compatibility; compatibility is not verification; an artifact is not reliance; payment is not evidence quality.
4. **Make limitations first-class.** Missing evidence, uncertainty, restrictions, supersession, and bounded purpose

travel with the artifact.

5. **Country rules live above the core.** Australian rules are versioned determinations, not hard-coded changes to the global evidence envelope.
6. **Commercial work must produce reusable assets.** Every paid engagement must improve the common plane.
7. **External authority remains external.** Verifiers, buyers, registries, regulators, and other institutions own their determinations.

Source basis: S4, S6, S7, S8..

5 Customer Problem

Agricultural technology companies already create extensive operational data: livestock identity and movement records; intervention and dosing events; water-flow and injection telemetry; product-batch and assay records; sensor readings and calibration histories; trial documents; remote-sensing outputs; processor measurements; model inputs and outputs; and payment or settlement events.

These records are valuable for operating a product or service, but they are not automatically suitable for institutional reliance. A buyer, auditor, verifier, registry, lender, research sponsor, or reporting entity must usually answer a larger set of questions:

- Where did the record originate, and which system or person had authority to create it?
- Can the record be linked to the correct animal, herd, property, intervention, batch, facility, model, or project?
- Was the record changed, replayed, duplicated, superseded, or rejected?
- Which schema, software, model, method, and version interpreted it?
- What evidence was missing, inconsistent, or outside the model's supported domain?
- Who reviewed the result, under which role and authority, and with what rationale?
- Which policy, method, registry, buyer, assurance, or contract profile applied on that date?
- Who held the rights to the underlying asset, data, and environmental claim?
- How were benefits allocated, capped, adjusted, retired, reversed, or prevented from duplicate claiming?
- What did the external institution actually accept, reject, or qualify, for which purpose and within which limits?

Most agricultural technology companies do not want to build bespoke provenance rails, canonical evidence graphs, immutable review histories, methodology adapters, rights and allocation engines, auditor-facing artifact workflows, and continuing policy operations. Even when they can build parts of that stack, the work competes with product delivery and remains difficult to standardize across buyers and jurisdictions.

AgEvidence's opportunity is to fill that gap without displacing the operational system. The customer continues to build and operate its product. AgEvidence makes the resulting evidence inspectable, portable, reviewable, and usable by the institution that must rely on it.

5.1 Current alternatives and their limits

Alternative	What it does well	Structural limitation	AgEvidence response
Customer builds internally	Maximum control and product-specific knowledge.	High opportunity cost; repeated reinvention; weak portability; ongoing policy and reviewer burden.	Reusable evidence tract and managed operations.

Carbon-accounting platform	Consolidated calculations and reports.	Often downstream of source provenance; may couple data to one calculation or reporting workflow.	Preserve source authority and separate evidence from method determinations.
Consultancy or audit preparation	Flexible interpretation and document assembly.	Manual reconstruction; limited repeatability; artifacts may depend on analyst knowledge.	Executable lineage, immutable review, portable artifact, and local verification.
Registry or verifier portal	Direct connection to a specific external process.	Institution-specific; may not preserve a reusable cross-institution evidence layer.	One evidence base, multiple bounded institutional profiles.
Farm or livestock platform extension	Close to operational data.	Customer-specific stack; may not support multi-party rights, allocation, policy portability, or external reliance.	Thin-waist integration without platform replacement.

Source basis: S2, S4, S7, S8..

6 Target Market and Initial Account Sequence

6.1 Initial market boundary

The first product market is deliberately constrained to:

- Australia and closely adjacent Australia-New Zealand livestock evidence flows where cross-border portability is itself part of the test;
- livestock and livestock-related technology;
- methane intervention, delivery, measurement, modeling, and verification evidence;
- processor and supply-chain evidence;
- externally reviewed agricultural climate claims; and
- Scope 3 and inseting programs in which evidence must move among producers, technology providers, project sponsors, buyers, verifiers, and reporting entities.

New geographies and non-livestock verticals should wait until the initial product demonstrates recurring revenue, external reliance, and reusable evidence components.

6.2 Primary customer groups

6.2.1 Event-first technology providers

These companies operate live telemetry, APIs, devices, workflow systems, or production platforms. They are the best fit for the signed Event Inbox path.

Representative needs: signed event mapping; source identity; device and software version continuity; calibration and exception records; idempotency and replay; source-to-artifact traceability; result webhooks; and production event bridges.

6.2.2 Source-record-first companies

Biological-input, research, laboratory, regulatory, and document-heavy organizations often maintain evidence in controlled files rather than continuous event streams.

Representative needs: source manifests; file hashes; authority and consent metadata; product-batch, assay, trial, safety, and efficacy provenance; version-controlled packages; and a path to external review without copying every file into a central application.

6.2.3 Measurement and processor platforms

These companies connect animal identity, facility measurement, grading, model interpretation, settlement, and buyer decisions.

Representative needs: model-version manifests; facility and device identity; processor evidence chains; measurement-to-decision lineage; buyer and auditor artifacts; and settlement-related reliance records.

6.2.4 Corporate buyers and program operators

Processors, retailers, food companies, financial institutions, research programs, and supply-chain sponsors require evidence from upstream producers and technology providers.

Representative needs: supplier comparability; bounded Scope 3 evidence; claim-allocation controls; methodology portability; assurance-ready packages; limitations; and reliance documentation.

6.3 Wave 1 target sequence

The supplied account dossiers prioritize six intensive Wave 1 targets. These are internal hypotheses, not proof of buying intent.

Target	Best entry	First technical action	Paid wedge hypothesis
DIT AgTech	Event Inbox	Map one intervention, delivery, dosing, or device event and preserve source/device lineage.	Livestock signed-event architecture and production bridge.
MEQ Solutions	Event Inbox	Map one measurement or model-completed event with facility, device/-model version, and buyer state.	Processor evidence architecture and enterprise buyer/auditor artifact.
Agscent	Source record + model review	Submit one measurement report and model reference; inspect calibration, cohort, and provenance gaps.	Verification Readiness Cycle.
Cibo Labs	Source record + model run	Submit one spatial model output reference with temporal and source commitments.	Remote-sensing evidence architecture, gated by reuse.
Rumin8	Source record + API	Submit one controlled trial, product, or regulatory reference and map batch/intervention identity.	Protocol Evidence Implementation.
Sea Forest	Source record	Submit one feed-additive evidence reference and map supply, batch, feeding, cohort, rights, and claims lineage.	Scope 3 reliance architecture.

6.4 Target-account qualification

No target becomes commercially qualified merely because it is funded, prominent, technically sophisticated, or strategically adjacent. Qualification requires all four of the following:

1. a meaningful technical action using a representative event, source manifest, or controlled fixture;
2. a real evidence, integration, production, assurance, or review gap;
3. a named external milestone, obligation, reviewer, buyer, sponsor, or relying institution; and
4. a paid scope that creates reusable capability or a reusable artifact profile.

Before quoting, the team must validate the technical sponsor, economic buyer, contracting authority, available project budget, restrictions on capital, external deadline, credible alternative cost, data rights, security requirements, acceptance criteria, and the customer's willingness to pay.

Source basis: S2, S3, S4, S5, S7, S8..

7 Customer Jobs to Be Done

The product should be designed around six principal customer jobs. These jobs are more stable than any individual feature list and should govern prioritization during the first 12 months.

7.1 Job 1: Connect evidence without replacing the source platform

A customer must be able to map a small number of material operational events or controlled files into AgEvidence while preserving its existing application, device, database, file store, user workflow, and source-system authority.

Success condition: the first integration can be delivered as a one-way evidence bridge, controlled source reference, or narrow API mapping rather than a core-system migration.

7.2 Job 2: Establish provenance and integrity

The customer must be able to show what was observed, when, by which system or person, under whose authority, using which schema and software version, and whether the record remains intact.

Success condition: a reviewer can traverse from artifact to projection to source commitment and verify integrity without privileged access to undocumented internal steps.

7.3 Job 3: Identify evidence gaps before formal review

Fixture-backed processing and structured review should expose missing fields, broken identity chains, absent calibration records, unclear authority, unsupported assumptions, model-domain limits, and insufficient review evidence before a costly audit or buyer review.

Success condition: the system produces a bounded gap report with owner, severity, remediation path, and relevance to the named external obligation.

7.4 Job 4: Evaluate the same evidence under multiple policies

The underlying evidence must remain stable while method, country, buyer, registry, assurance, or contract requirements are applied through versioned profiles.

Success condition: a new policy profile can generate a new compatibility determination and artifact without rewriting historical evidence.

7.5 Job 5: Assemble bounded reliance artifacts

The system must create a portable package for a named institution, decision, and purpose, including evidence lineage, review history, method/profile version, limitations, verification instructions, and explicit non-claims.

Success condition: the artifact states exactly what it supports and does not imply broader approval.

7.6 Job 6: Operate recurring evidence flows

After the initial architecture is proven, AgEvidence should manage ingestion, exceptions, schema evolution, policy updates, review queues, artifact regeneration, and reliance records as a recurring service.

Success condition: recurring operations reduce customer reconstruction effort and AgEvidence's bespoke engineering effort per artifact.

Source basis: S4, S7, S8..

8 Product Architecture Overview

8.1 Architecture thesis


The product is a layered evidence system, not a single ledger and not a single application. It must preserve three distinct functional tiers:

Tier	Owns	Must not own	Primary output
------	------	--------------	----------------

Global evidence layer	Canonical envelope, stable event semantics, source commitments, identity and rights bindings, parent links, signatures, limitations, append-only lineage, bundle manifests.	Local eligibility, final scientific judgment, registry acceptance, universal claim ownership, or government authority.	Evidence graph, receipts, manifests, verifier inputs.
Country and method layer	Versioned applicability, baseline, method, model, uncertainty, assurance, data, and claim profiles.	Rewriting evidence history; cryptographic implementation; final external authority.	Compatibility determination, gap list, method/country artifact profile.
Institutional reliance layer	Buyer, verifier, registry, sponsor, lender, processor, contractual, or reporting decision.	Rewriting evidence; representing one decision as universal acceptance.	Acceptance, rejection, qualification, reliance artifact, allocation/-payment/retirement records.

8.2 Ingestion path A: signed operational events

Mature operational systems submit versioned JSON events into the Evidence Event Inbox.



The working materials describe signed intake, idempotency, append-only persistence, evidence projection, receipt requests, replay, webhooks, operations visibility, and dead-letter handling. Those capabilities are the commercial spine, but production claims remain gated by security, load, failure-recovery, data-rights, and external review requirements.

8.3 Ingestion path B: controlled source records

Document-heavy customers retain raw files in controlled infrastructure and submit a manifest containing:

- file hash or cryptographic commitment;
- controlled location or reference;
- source type and format;
- authority, consent, and data-use metadata;
- creation date, effective date, and version;
- project, intervention, batch, cohort, facility, or model identifiers;
- retention, access, and confidentiality classification; and
- limitations or known data-quality conditions.

This path is appropriate for laboratory reports, batch certificates, trial spreadsheets, toxicology records, regulatory documents, spatial outputs, and other evidence that should not be copied indiscriminately into a central application. It also provides the fastest evaluation path for intervention companies before they implement a production event stream.

8.4 Evidence processing lifecycle

1. Create a developer project and identify the named external obligation.
2. Register source systems, identities, roles, authority, and data rights.

3. Ingest signed events or controlled source manifests.
4. Validate schema, identity, idempotency, signature, and required metadata.
5. Project accepted records into append-only evidence structures.
6. Run fixture-backed evidence evaluation and candidate generation.
7. Surface candidate evidence, missing evidence, exceptions, and unsupported assumptions.
8. Record append-only human review decisions with authority and rationale.
9. Apply a versioned method, country, buyer, assurance, or institutional profile.
10. Produce a scoped quote or artifact order tied to exact evidence and deliverables.
11. Assemble the reliance artifact with manifest, lineage, limitations, and verification instructions.
12. Verify the package locally and record the result.
13. Record any external acceptance, rejection, qualification, or reliance event separately.
14. Operate recurring ingestion, exceptions, policy updates, artifact regeneration, and support.

Source basis: S1, S4, S5, S6, S7, S8..

9 Canonical Graph Model

A single evidence graph is insufficient for multi-party agricultural claims. Evidence validity, legal rights, allocation, payment, and institutional acceptance change independently and must remain separately inspectable.

Graph	Core question	Representative contents
Evidence	What occurred, according to which sources?	Events, source records, intervention and batch identity, delivery and dosing, measurements, calibration, model inputs and outputs, uncertainty, signatures, timestamps, limitations, and parent links.
Rights	Who had legal or contractual authority over the asset, data, and participation?	Land, lease, herd, asset, and facility authority; producer consent; data-use permissions; delegated rights; project participation; assignment and restrictions.
Allocation	How is the recognized environmental outcome divided and controlled?	Co-funding shares, supply-chain boundaries, allocation formulas, claim caps, unallocated balance, buffers, reversals, adjustments, retirement, and duplicate-claim controls.
Payment	What financial events occurred and what is their status?	Contributions, invoices, incentives, fees, revenue sharing, escrow, producer payment, settlement, reversal, and dispute events.
Reliance	Which institution accepted, rejected, or qualified which evidence for which purpose?	Verifier findings, buyer decisions, registry actions, reporting references, research decisions, lender or insurer decisions, requested evidence, limitations, and effective dates.

9.1 Graph invariants

1. The evidence root does not change merely because a method, policy, buyer requirement, or contract changes.
2. Rights, allocation, payment, and reliance events extend the history; they do not retroactively rewrite observations.
3. Financial contribution does not automatically determine claim ownership.
4. A payment event does not imply evidence quality, verification, or reliance.
5. One institution's reliance does not imply universal acceptance.

6. A rejected or qualified reliance event remains part of the record and can be as valuable for product learning as an acceptance.
7. Every artifact states which graph versions and policy/profile versions it contains.

9.2 Minimum entity relationships for the first livestock wedge

The first production model should support, at minimum:

- organization, person, role, and authority;
- developer project, customer project, and external obligation;
- source system, source record, external object, and event;
- animal, herd, cohort, property, facility, intervention, product, and batch identifiers where applicable;
- measurement device, calibration record, model, model version, and model run;
- evidence projection, candidate, gap, reviewer decision, and limitation;
- method/country/institutional profile and compatibility determination;
- artifact engagement, artifact order, package, manifest, and verification result;
- rights, allocation, payment, and reliance events; and
- repository release, schema version, handler version, and artifact-profile version used to deliver the engagement.

Source basis: S6, S7, S8..

10 Trust-State Model

The product must prevent **validity laundering**: treating one valid property as proof of every other property. This is not a disclosure footnote; it is a core data-model and product-behavior requirement.

State family	Meaning	Examples
Cryptographic validity	Hash, signature, canonicalization, parent-link, and append-only sequence integrity.	valid, invalid, unverifiable
Scientific or measurement validity	Calibration, sampling, uncertainty, empirical appropriateness, model domain, and data quality.	supported, qualified, insufficient
Method compatibility	Whether the evidence satisfies the requirements of a specific method, protocol, or project profile.	compatible, partial, incompatible
Human review status	Explicit domain-expert acceptance, rejection, qualification, rationale, and escalation.	pending, accepted, rejected
Verification determination	Finding made by an appropriately authorized external verifier or assurance provider.	verified, qualified, adverse, none
Artifact assembly status	Structural state of package compilation, validation, signing, and export.	draft, assembled, failed, superseded
Institutional reliance	Whether a named institution used the artifact for a defined decision and within defined limits.	accepted, rejected, qualified, requested-more
Claim-allocation status	Ownership, co-funding, boundary, caps, retirement, reversal, and duplicate-claim controls.	allocated, partially allocated, disputed
Payment and settlement status	Commercial flows separate from evidence quality and reliance.	proposed, invoiced, collected, settled
Revenue status	Internal accounting stage, also separate from product and reliance states.	sandbox, proposed, contracted, invoiced, collected, recognized

10.1 Required state-transition controls

- A cryptographically valid receipt cannot automatically advance scientific, method, verification, reliance, allocation, payment, or revenue states.
- Model output must enter a review-required state unless an explicitly approved workflow says otherwise; the model does not become institutional authority.
- Artifact assembly may occur before external reliance, but the artifact must remain labeled as unrelayed upon until a separate event is recorded.
- Sandbox quotes and orders must never populate contracted, invoiced, collected, or recognized revenue fields.
- A country adapter may issue a compatibility result, not government approval.
- Any supersession must preserve prior determinations, effective dates, and the reason for change.
- User interfaces, API responses, exports, contracts, and dashboards must display the relevant state family rather than using an ambiguous global “valid” status.

Prohibited inference examples

- “Signature verified” may not be rendered as “verified emissions reduction.”
- “Method-compatible” may not be rendered as “approved by the regulator.”
- “Artifact assembled” may not be rendered as “accepted by the buyer.”
- “Buyer relied” may not be rendered as “universally accepted.”
- “Customer paid” may not be rendered as “scientifically validated.”

Source basis: S1, S4, S5, S6, S7, S8..

11 Core Product Modules

The product catalog below distinguishes the desired operating product from the current repository scaffold. Each module requires a capability owner, acceptance tests, explicit limitations, and a decision about whether it belongs in the open commons or paid commercial layer.

11.1 Evidence Event Inbox

Purpose: receive a narrow vocabulary of versioned operational events from existing systems and preserve source authority.

Required capabilities:

- versioned event API and schema registry;
- source identity, authentication, signing, and key-management boundary;
- idempotency, duplicate detection, accepted/rejected state, and immutable raw intake;
- external-object mapping and evidence projection;
- replay, dead-letter handling, retry policy, outbox, webhook delivery, and operations inspection;
- correlation identifiers across source event, projection, receipt, artifact, and result; and
- metrics for acceptance, latency, failure, replay, support, and release attribution.

11.2 Source Manifest Manager

Purpose: preserve controlled references to documents and datasets without centralizing unnecessary raw content.

Required capabilities:

- file hashes or commitments; controlled location; source type; date and version;
- authority, consent, data-use, confidentiality, retention, and access metadata;
- source-project, intervention, batch, cohort, model, or facility links;
- manifest validation, supersession, availability, and access-failure states;
- controlled disclosure and de-identification rules; and
- evidence that a source existed in a defined state at a defined time.

11.3 Evidence Projection Engine

Purpose: map heterogeneous source objects into canonical evidence records while preserving lineage.

Required capabilities:

- mapping definitions with version, owner, effective date, and test fixtures;
- identity correlation and conflict handling;
- event ordering, causality, parent links, and intervention reconstruction;
- unit normalization without loss of original values;
- source-to-projection lineage and reversible inspection;
- exception and unsupported-field handling; and
- explicit classing of mappings as supported, source-record supported, profile extension, new-event candidate, or deferred domain family.

11.4 Evidence Gap Analyzer

Purpose: identify evidence insufficiency before formal external review.

Required capabilities:

- fixture-backed evaluation and reproducible model runs;
- missing-field, identity-chain, calibration, authority, and temporal-gap detection;
- unsupported assumption and model-domain warnings;
- method-specific gap profiles;
- severity, owner, remediation, external-obligation relevance, and disposition;
- reusable gap taxonomy and aggregation across customers; and
- separation between machine-generated candidate gaps and human-approved findings.

11.5 Human Review Ledger

Purpose: make domain judgment explicit, append-only, attributable, and reviewable.

Required capabilities:

- reviewer identity, role, authority, conflict status, and effective scope;
- acceptance, rejection, qualification, request-for-evidence, and escalation;
- rationale, supporting references, limitations, and linked evidence version;
- immutable decision history with supersession rather than destructive editing;
- separation of customer review, AgEvidence review, scientific review, and external verifier determination; and

- exportable review trail for artifacts.

11.6 Policy and Method Adapter Layer

Purpose: evaluate preserved evidence under versioned country, method, assurance, buyer, registry, and contract requirements.

Required capabilities:

- profile identity, version, owner, source basis, effective date, jurisdiction, and applicability;
- rule declarations for required evidence, calculations, uncertainty, eligibility, review, and limitations;
- side-by-side evaluation under multiple profiles;
- migration, supersession, deprecation, and historical replay;
- explicit unsupported and externally dependent rules; and
- no hard-coded local rules in the cryptographic or global evidence core.

11.7 Claim and Rights Engine

Purpose: preserve authority and govern multi-party allocation without implying universal legal rules.

Required capabilities:

- participant, asset, data, project, and claim rights;
- consent, delegation, assignment, restriction, and expiration;
- co-funding, allocation formulas, caps, buffers, duplicate-claim controls, retirement, reversal, and dispute;
- audit trail and effective dates;
- explicit legal and contractual source references; and
- controls preventing payment contribution from being treated as automatic claim ownership.

11.8 Reliance Artifact Builder

Purpose: assemble a portable package for a named institution and decision.

Required capabilities:

- package manifest, index, evidence lineage, rights and allocation summary, review history, profile version, limitations, and non-claims;
- local verification instructions and machine-readable verification metadata;
- institution-specific artifact profile and disclosure boundary;
- deterministic assembly from versioned inputs;
- export to agreed portable formats; and
- supersession and revocation record.

11.9 Reliance Registry

Purpose: record what an external institution actually did, separately from evidence and artifact states.

Required capabilities:

- named institution and authorized decision role;
- decision type, date, purpose, scope, evidence and artifact version;

- acceptance, rejection, qualification, request for more evidence, or no decision;
- reliance limits, duration, conditions, and supporting documentation;
- relation to allocation, payment, disclosure, or retirement events; and
- external-authority provenance.

11.10 Managed Evidence Operations

Purpose: convert one-time architecture into recurring operational value.

Required capabilities:

- recurring source monitoring, queue and webhook operations, replay, dead-letter resolution, and incident management;
- schema evolution, profile updates, migration testing, and release attribution;
- review-queue operation, artifact regeneration, and evidence-gap reporting;
- audit, buyer, verifier, and customer support;
- service reporting, service levels, change management, and access governance; and
- continuous reduction in bespoke code and analyst reconstruction effort.

Source basis: S1, S4, S5, S6, S7, S8..

12 Capability Register: Implemented, Partial, and Absent

The launch must begin with an evidence-based capability register. The table below is a proposal-level synthesis of the supplied repository descriptions; the technical owner must verify each line against the selected release before external use.

Capability	Working status	Launch implication
Signed integration event endpoint	Implemented scaffold	Use as the primary event-first wedge after clean-machine validation and security review.
Append-only integration event persistence	Implemented scaffold	Preserve original intake and rejected/accepted state; verify immutability semantics.
External object mapping and evidence projection	Implemented scaffold	Use for first customer mapping; document lineage and mapping version.
Receipt request/outbox projection	Implemented scaffold	Demonstrate, but do not infer production trust or institutional acceptance.
Replay, webhooks, result events, dead-letter handling	Implemented scaffold	Instrument and test failure recovery, retries, idempotency, and operations burden.
Controlled source-record path	Implemented scaffold	Use for trials, labs, regulatory documents, and other source-record-first targets.
Fixture-backed model runs and candidate/gap review	Implemented scaffold	Keep model output review-required; validate fixture scope and limitations.
Append-only human review decisions	Implemented scaffold	Verify role/authority controls and exportable review history.
Sandbox pricing quote and artifact order path	Implemented scaffold	Demonstrate product logic only; keep separate from approved proposals, contracts, payments, and revenue.
Artifact assembly metadata and local verification path	Implemented scaffold	Validate deterministic assembly, package completeness, and independent clean-machine verification.
Self-service guide and role-orientation pack	Implemented scaffold	Treat as product surface; validate every guide against the tagged release.

Production security architecture	Partial / requires hardening	Requires threat model, key management, secrets, access control, logging, vulnerability process, and customer-specific review.
Production load, availability, and recovery	Partial / requires hardening	Establish SLOs, capacity tests, backup/recovery, queue and webhook behavior, and incident procedures.
Complete Australian method and country profiles	Partial / requires hardening	Build only profiles required by named paid obligations; do not claim authority.
Complete rights, allocation, and duplicate-claim controls	Partial / requires hardening	Prioritize when a paid co-claiming or Scope 3 use case requires them.
Independent verifier or buyer acceptance	Absent / externally dependent	Must be earned and recorded as a separate external reliance event.
Regulatory or registry approval	Absent / externally dependent	Remains with the competent external authority; product may support evidence and compatibility only.
Scientific efficacy of customer interventions	Absent / externally dependent	Must come from approved scientific evidence and authorized reviewers, not repository behavior.
Recognized revenue and recurring managed customers	Absent / externally dependent	Must be tracked separately from sandbox, proposals, contracts, invoices, and cash.

12.1 Release truth requirements

Before private activation, the technical owner must produce:

1. the exact repository SHA and signed or otherwise controlled release tag;
2. test results for Rails, Rust, Python, integrations, replay, webhooks, and local verification;
3. clean-machine installation and end-to-end demo evidence;
4. an inventory of synthetic, fixture, controlled, and real-data paths;
5. a known-limitations register;
6. a schema and event vocabulary manifest;
7. a dependency and trust-boundary diagram;
8. a production-hardening backlog linked to customer obligations rather than speculative completeness; and
9. a rule that conflicting orientation or campaign documents are corrected when they diverge from canonical implementation specifications.

Source basis: S1, S4, S5, S8..

13 Data Contract and Reliance Artifact Anatomy

13.1 Minimum evidence envelope

The exact schema remains a technical design artifact, but the proposal requires every evidence envelope or source projection to carry enough information to preserve provenance and bounded interpretation.

Field family	Purpose	Minimum requirement
Identity	Correlate evidence to the correct entity.	Event/source ID, project ID, source-system ID, object/entity IDs, parent links.
Authority	Establish who may create or provide the evidence.	Organization/person/role, delegated authority, consent or rights reference.

Time	Preserve occurrence, observation, ingestion, and effective time.	Event time, observed time, received time, time-zone, sequence where required.
Schema and software	Reproduce interpretation.	Event/schema version, producer software version, mapping version, handler version.
Content commitment	Preserve integrity without unnecessary copying.	Canonical payload or file hash, signature meta-data, original units and values.
Context	Make the evidence scientifically and operationally interpretable.	Intervention, batch, cohort, property/facility, device, calibration, model, method references as applicable.
Limitations	Prevent overclaiming.	Missing fields, uncertainty, unsupported assumptions, access limits, confidentiality, model-domain restrictions.
Lineage	Trace transformations and decisions.	Source-to-projection links, model run, candidate, review, profile, artifact, verification, reliance links.

13.2 Reliance artifact structure

Every external-facing artifact should contain, at minimum:

1. **Cover and purpose.** Named institution, decision context, customer, project, artifact ID, version, date, and confidentiality.
2. **Executive reliance statement.** What the package supports, what it does not support, and the exact question presented.
3. **Manifest.** Files, receipts, evidence roots, schema/profile versions, software release, and verification instructions.
4. **Source and authority register.** Source systems, controlled references, roles, rights, and permissions.
5. **Evidence lineage.** Event/source-to-projection-to-review-to-artifact trace.
6. **Gap and limitation register.** Missing evidence, qualifications, uncertainty, unsupported assumptions, and open issues.
7. **Human review history.** Reviewer identity, authority, decision, rationale, and supporting references.
8. **Method or institutional profile.** Version, applicability, effective date, compatibility determination, and externally dependent requirements.
9. **Rights and allocation summary.** Only when relevant and supported; with explicit legal/contractual basis.
10. **Verification result.** Local structural and cryptographic verification, kept separate from external verifier determination.
11. **Reliance record.** Added only when the named institution records acceptance, rejection, qualification, or request for more evidence.
12. **Supersession and revocation.** How future versions relate to this package.

13.3 Artifact acceptance criteria

An artifact is structurally complete only when:

- ✓ every included object appears in the manifest and can be located;
- ✓ every projection links to a source event or source commitment;
- ✓ every human decision identifies the evidence version and reviewer authority;
- ✓ every method or institutional result identifies its profile version and effective date;
- ✓ limitations and missing evidence are visible, not buried in implementation logs;

- ✓ local verification succeeds on a clean machine using published instructions;
- ✓ the package contains no unsupported implication of scientific, regulatory, verifier, registry, buyer, or government approval; and
- ✓ the external reliance state remains unset until an external decision is separately recorded.

Source basis: S1, S4, S6, S7, S8..

14 Security, Privacy, Rights, and Operations Requirements

Production hardening must be driven by the first contracted source and relying-party obligation. The proposal nevertheless establishes minimum control domains that cannot be deferred indefinitely.

14.1 Security requirements

- threat model covering source authentication, signing, replay, impersonation, tampering, data exfiltration, privilege escalation, webhook abuse, dependency compromise, and artifact substitution;
- documented separation among integration authentication, receipt trust, application authorization, model execution, review authority, and institutional reliance;
- key-management lifecycle, rotation, revocation, secure storage, and incident response;
- role-based access control and least privilege across customer, AgEvidence, reviewer, operations, and external-institution roles;
- encryption in transit and at rest where appropriate, with customer-specific requirements documented;
- audit logging, monitoring, vulnerability management, dependency inventory, and release provenance;
- rate limiting, idempotency, replay controls, webhook signing, retry policy, and dead-letter security; and
- independent review before representing the system as production-ready.

14.2 Privacy and data stewardship

- collect and centralize only data necessary for the bounded evidence purpose;
- prefer controlled source references for sensitive, large, or rights-restricted files;
- classify personal, commercial, scientific, producer, location, and contractual data;
- define access, retention, deletion, legal hold, de-identification, and disclosure rules;
- document cross-border storage or processing where relevant;
- ensure customer and producer authority is represented in the rights graph; and
- obtain data-steward approval before publication, benchmarking, or reuse.

14.3 Operational requirements

- observable acceptance, rejection, latency, queue, webhook, replay, and dead-letter metrics;
- runbooks for source outage, malformed event, duplicate event, signature failure, projection failure, receipt failure, webhook failure, artifact failure, and verifier question;
- release attribution connecting repository SHA, schema, handler, mapping, profile, artifact, external decision, and commercial state;
- backup, recovery, migration, and supersession tests;
- service-level objectives introduced only after measured customer requirements;

- support burden tracked per activation and per artifact; and
- preservation of original records and failure evidence during incident remediation.

Non-negotiable guardrail

Security hardening is not a generic precondition for all customer discovery, but production use is never implied by the existence of a working scaffold. Each paid production proposal must state the current control posture, customer-specific requirements, work required, residual risk, and acceptance owner.

Source basis: S1, S4, S6, S8..

15 Product Packaging

The initial catalog should function as an internal commercial architecture, not a public fixed-price list. Each offer must be tied to a named evidence obligation, exact deliverables, customer data boundary, repository release, acceptance criteria, and a reuse hypothesis.

Offer	Purpose	Indicative deliverables	Entry anchor
Evidence Architecture Sprint	First paid diagnostic and design engagement.	Source inventory; event/source mapping; authority and rights map; canonical evidence graph; gap map; named relying party; implementation plan; reusable as-set plan.	US\$25,000
Verification Readiness Cycle	Prepare for a defined external review.	Evidence sufficiency assessment; review workflow; method/profile mapping; remediation plan; reconstruction test; limitations register.	US\$60,000
Enterprise Reliance Artifact	Assemble a bounded package for a named institution.	Portable artifact; manifest; evidence and review lineage; profile; limitations; local verification instructions; reliance-question statement.	US\$30,000
Protocol Evidence Implementation	Implement production evidence flows.	Production bridge; schemas and handlers; method/country profile; controls; tests; artifact profile; operational handoff.	US\$100,000
Managed Evidence Plane	Operate recurring evidence and reliance flows.	Managed ingestion; exceptions; review operations; policy/profile updates; artifact regeneration; support; reliance registry; service reporting.	US\$150,000 annually

These figures are internal planning assumptions from the supplied campaign materials. They are not validated market prices, quotes, commitments, forecasts, or representations of willingness to pay.

15.1 Packaging refinement

Before a formal quote, the commercial owner must distinguish:

- campaign entry anchor versus account-specific bounded price;
- architecture/design work versus production hardening;
- standard reusable deliverables versus customer-specific work;
- included external-review interaction versus third-party verifier or legal costs;
- one-time implementation versus recurring managed operations;
- open assets versus confidential customer configuration and controlled data; and

- successful artifact assembly versus any promise of external acceptance.

The commercial diligence documents model materially higher account-specific ranges for some targets. Those ranges reflect project value, complexity, external obligation, and risk; they should not be presented alongside the entry anchors as equivalent price lists. The correct relationship is:

Entry anchor → qualified evidence obligation → bounded value model → approved formal quote

15.2 Standard deliverable definitions

Deliverable	Definition of done
Source inventory	Enumerates authoritative sources, owners, access, rights, update frequency, versions, and known limitations.
Event/source mapping	Maps representative objects into the current contract; identifies profile extensions, new-event candidates, and deferred domain families.
Evidence graph	Reconstructs the relevant project or intervention with source lineage and explicit missing evidence.
Gap map	Prioritized findings tied to the named external obligation, with owner and remediation path.
Profile mapping	Versioned interpretation of method, country, assurance, buyer, or institutional requirements; external authority clearly separated.
Reliance artifact	Deterministic package containing manifest, lineage, review, profile, limitations, verification instructions, and bounded purpose.
Production bridge	Tested integration with operational monitoring, failure handling, access controls, release attribution, and handoff.
Managed operations	Recurring runbook, service reporting, exception management, release/profile updates, artifact regeneration, and support boundary.

Source basis: S2, S3, S4, S8..

16 Pricing Methodology and Commercial Controls

16.1 Value-based fee ceiling

AgEvidence should price against the incremental value and risk reduction of a named project, not against seats or technical volume.

For each engagement, define:

V_p = incremental value of the underlying customer project,
 C_a = credible cost of the customer's alternative,
 B = available and authorized project budget,
 R = AgEvidence delivery capacity and acceptable risk,
 α = maximum defensible value-capture fraction.

A defensible fee ceiling is:

$$F_{\max} = \min(\alpha V_p, C_a, B, R)$$

The final price should also reflect source-system count and complexity; data quality; identity and rights complexity; method and policy uncertainty; external deadline; production-hardening requirements; reviewer interaction; artifact scope; evidence reuse potential; and recurring managed-service opportunity.

16.2 Four independent commercial caps

No single value model is sufficient. A quote must remain below four independent caps:

1. **Value cap:** a defensible fraction of incremental project value or avoided risk.
2. **Alternative-cost cap:** the credible cost of internal build, consultancy, audit reconstruction, delayed deployment, or other substitute.
3. **Budget cap:** the customer's authorized, unrestricted, and procurable budget.
4. **Delivery-risk cap:** AgEvidence's capacity to deliver without disproportionate security, scientific, legal, operational, or reputational risk.

16.3 Commercial state discipline

The operating ledger must track these values separately:

Sandbox quoted $\neq$ proposed $\neq$ contracted $\neq$ invoiced $\neq$ collected $\neq$ recognized

- Sandbox quotes demonstrate product logic only.
- Formal proposals must use an approved commercial process outside the sandbox.
- Contracts must identify deliverables, payment schedule, start date, data boundary, acceptance owner, and change control.
- Cash collection does not establish scientific validity or institutional reliance.
- Recognized revenue follows AgEvidence's accounting policy, not artifact or payment status alone.

16.4 Quote approval packet

Every quote should include:

1. named customer project and external obligation;
2. source-system or source-record scope;
3. current repository capability and known gaps;
4. deliverables and excluded work;
5. acceptance tests and acceptance authority;
6. data, rights, security, privacy, and retention boundary;
7. external-review plan and non-guarantee of acceptance;
8. reusable assets expected from the work;
9. price logic and the four independent caps;
10. payment schedule and revenue-state treatment;
11. production-hardening assumptions; and
12. stop, change-order, and dispute conditions.

Source basis: S2, S3, S4, S8..

17 Go-to-Market Model

17.1 Zero-sales-team activation


The initial motion should be developer-led and evidence-led rather than sales-led. It begins with a reproducible product surface, not a generic demo funnel.

The minimum campaign package is:

- a frozen reference release and reproducibility manifest;
- open schemas, event examples, source-manifest examples, and synthetic livestock fixtures;
- an external-safe self-service guide index;
- clean-machine setup and local verification instructions;
- Python and TypeScript examples where supported;
- event-inbox, source-record, model-review, pricing/artifact, webhook/operations, and troubleshooting guides;
- a technical preprint with strictly bounded contribution claims;
- role-specific internal orientation for executive, GTM, engineering, integrations, trust, model, country-policy, and operations owners; and
- targeted technical outreach that asks each company to perform one specific evidence action.

The target should not be asked for a general discovery call or platform demo as the first step. It should receive a relevant guide and a narrow challenge, such as mapping one intervention event, submitting one trial manifest, or adding one measurement/model record.

17.2 Conversion sequence



17.3 Outreach proposition

Build and operate your product. AgEvidence will help make its evidence inspectable, portable, and usable by the institution that must rely on it.

The proposition is intentionally operational. It does not begin with carbon-market abstractions or a claim that AgEvidence replaces existing systems. It begins with the customer's evidence obligation and the smallest compatible source action.

17.4 Private-before-public sequence

1. Validate the guide pack and release internally on clean machines.
2. Activate five to eight private targets using controlled or synthetic paths.
3. Obtain at least three completed technical actions and two real evidence mappings.
4. Convert qualified gaps into one to three proposals.
5. Make public release only after critical guide blockers are resolved and claims language is approved.
6. Keep account dossiers, commercial controls, revenue gates, and internal assumptions out of the external self-service path.

17.5 Sales boundary

No broad enterprise funnel, quota-carrying sales team, broad paid media program, or complex CRM process should be built until the team proves:

- repeatable technical activation;
- a named evidence gap and external obligation;
- willingness to pay for the bounded engagement;
- a reusable implementation pattern;
- a real external review or reliance pathway; and
- economics that improve with reuse rather than require expanding bespoke services.

Source basis: S4, S5, S8..

18 Open-Core and Commercial Boundary

18.1 Open evidence commons

AgEvidence should publish the minimum common layer required for compatibility, evaluation, reproducibility, and scientific scrutiny:

- canonical vocabulary and evidence-envelope specifications;
- schemas and sample payloads;
- synthetic and permissioned example fixtures;
- validation logic and conformance tests;
- reference manifests and local verification tooling;
- interoperability and reproducibility procedures;
- selected country and method profiles where rights and governance permit;
- benchmark definitions and non-proprietary reference datasets where rights permit; and
- scientific papers, technical preprints, and claims-and-limitations templates.

18.2 Paid commercial layer

AgEvidence should retain commercial value in:

- customer-specific mappings and production integrations;
- controlled datasets and confidential customer configurations;
- operational hardening, security, observability, and service levels;
- rights, allocation, and customer-specific contract logic;
- policy/profile maintenance and migration operations;
- human review workflows and reviewer operations;
- institution-specific artifact assembly and support;
- managed ingestion, exception handling, artifact regeneration, and reliance tracking;
- support, assurance coordination, and change management; and
- portfolio-wide sponsor programs and private deployment patterns.

18.3 Boundary test

An asset should generally be open when it lowers ecosystem integration cost, makes claims inspectable, permits independent conformance, or establishes a common vocabulary without exposing customer rights or confidential implementation.

An asset should generally remain commercial when it requires customer-specific data, ongoing operations, contractual authority, controlled configuration, service guarantees, external coordination, production liability, or recurring policy and artifact maintenance.

Asset	Open rationale	Commercial rationale	Proposed placement
Canonical vocabulary	Enables shared interpretation and publication.	Limited direct capture value.	Open
Core schemas and synthetic fixtures	Reduces integration friction and permits testing.	Customer extensions may be confidential.	Core open; extensions controlled
Local verifier	Builds trust and artifact portability.	Managed verification operations create service value.	Tool open; operations paid
Customer mapping	Not generally useful without context and rights.	High customer-specific labor and reuse value.	Paid
Country/method profile	Some profiles support public conformance.	Maintenance, proprietary interpretation, and institutional profiles create recurring value.	Selective
Artifact template	Common structure improves reviewer expectations.	Institution-specific assembly and support are costly.	Base open; assembly paid
Managed operations	No meaningful open substitute for service accountability.	Recurring operational value and risk.	Paid

18.4 Licensing and governance decisions

Before external release, the executive team must approve:

- the license for schemas, reference code, fixtures, verifier, and documentation;
- contributor and inbound-license policy;
- trademark and compatibility-mark rules;
- conformance-test governance, versioning, revocation, and supersession;
- customer-data and derivative-work boundaries;
- publication and dataset rights; and
- controls preventing open compatibility claims from being confused with regulatory, scientific, or institutional approval.

Source basis: S4, S6, S8..

19 Customer Qualification and Account Diligence

The zero-sales motion does not eliminate commercial diligence. It moves qualification closer to technical evidence and external obligations.

19.1 Five-dimensional diligence

Each account should be evaluated across at least five dimensions:

1. **Funding and runway mechanics.** Remaining unrestricted cash, program-specific capital, burn, procurement timing, and restrictions - not headline funding alone.
2. **Decision authority.** Technical sponsor, economic buyer, contracting authority, legal/security approvers, and procurement threshold.
3. **Evidence obligation.** Named buyer, verifier, registry, research sponsor, reporting deadline, project gate, financing event, or customer commitment.
4. **Technical fit.** Event-first or source-record-first path, source accessibility, identity and authority complexity, current contract fit, and hardening burden.
5. **Reuse and reference value.** Ability to create common livestock schemas, handlers, profiles, artifacts, and credible external learning.

19.2 Account evidence packet

Before a formal proposal, the account owner should prepare:

- dated company and product identity;
- public and customer-validated architecture hypothesis;
- evidence lifecycle and authoritative source list;
- current repository mapping class;
- missing capabilities and their reuse potential;
- named external obligation and decision date;
- technical sponsor, economic buyer, procurement and data-rights path;
- budget source and credible alternative cost;
- proposed first technical action;
- bounded paid scope and acceptance criteria; and
- claims that are verified, inferred, proposed, or prohibited.

19.3 Scoring model

The source dossiers use a prioritization score that rewards evidence structure, event fit, source-record fit, external obligation, accessibility, urgency, reuse, paid-wedge clarity, and reference value, while subtracting implementation burden. The exact weights may be revised, but the principle should remain: **priority is a function of evidence fit and external obligation, not company prestige alone.**

19.4 Disqualification conditions

An account should be deferred when:

- there is no accessible representative source or technical sponsor;
- no external obligation can be named;
- the customer wants a replacement operating platform;
- the requested domain requires a speculative new vocabulary with no reusable first wedge;
- budget or contracting authority cannot be established;
- data rights or security requirements cannot be met responsibly;
- the customer requires AgEvidence to claim scientific or regulatory authority it does not possess; or

- the work produces no reusable capability, artifact profile, or managed-operations path.

Source basis: S2, S3, S5, S8..

20 Ninety-Day Commercial Program

20.1 Program objective

The 90-day program should prove or disprove the bridge-first thesis with minimal irreversible investment. It must produce a commercial signal, an external review signal, and a reuse signal.

20.2 Day 0–30: proof phase

Timing	Workstream	Required output	Owner
Days 0–3	Authorization and canonicalization	Signed charter; livestock scope; decision rights; claims rule; currency; exact source and repository reconciliation plan.	Kendra / executive approver
Days 1–10	Executable surface	Tagged release; clean-machine validation; test register; event/source paths; synthetic fixtures; known limitations; guide corrections.	Technical owner
Days 4–12	Campaign package	External-safe guide index; role router; technical preprint alpha; claims matrix; target-specific activation notes.	Campaign / publication leads
Days 8–20	Private activation	Five to eight targets contacted; three completed technical actions; at least two real evidence mappings.	Commercial + technical
Days 12–23	Scope conversion	Gap and obligation analysis; reusable capability plan; one to three bounded formal proposals.	Commercial / integration
Days 24–30	Revenue gate	At least one paid engagement, or a documented continue/revise/stop decision based on evidence.	Kendra / executive approver

20.3 Day 31–60: first implementation and independent reconstruction

- complete the first paid Architecture Sprint or equivalent mapping;
- deliver one customer-specific fixture through the existing inbox/source path, projection, review, receipt/artifact, and verification sequence;
- implement only the minimum new livestock fields or event types required by the contract;
- instrument acceptance, latency, projection, outbox, webhook, replay, dead-letter, support, and release-attribution metrics;
- produce a bounded gap and production-hardening plan;
- identify the exact external reviewer or relying-party question; and
- commission an independent reconstruction or readiness review with documented exceptions and limitations.

20.4 Day 61–90: external artifact and reuse proof

- harden one source integration to the contracted standard;
- complete one method, country, buyer, or institutional profile needed by the engagement;
- assemble one external-review artifact and verify it independently;
- record structured reviewer feedback, including rejection or qualification where applicable;
- activate a second customer and measure reuse of schemas, handlers, projections, gap taxonomies, or artifact profile;

- create the first release-to-revenue-to-reliance lineage record; and
- present a 12-month continuation plan only if commercial, reuse, and reliance gates are met.

20.5 Day-30 gate

Continue only when the program has:

- ✓ one reproducible tagged release and zero unresolved critical guide blockers;
- ✓ at least three completed external technical actions;
- ✓ at least two representative evidence mappings;
- ✓ at least one named external obligation;
- ✓ one to three formal proposals; and
- ✓ at least one paid engagement or exceptionally strong documented evidence justifying a narrowly extended conversion period.

20.6 Day-90 gate

Continue into recurring productization only when the program has:

- ✓ one paid customer-specific evidence flow delivered against acceptance tests;
- ✓ one external artifact inspected by a named reviewer or institution;
- ✓ one reusable capability adopted by a second customer or second use case;
- ✓ a credible managed-operations requirement;
- ✓ a production-hardening plan grounded in measured customer needs; and
- ✓ no unresolved scientific, rights, security, legal, or claims issue that makes continuation irresponsible.

Source basis: S1, S4, S8..

21 Twelve-Month Productization Roadmap

The 12-month roadmap is not a commitment to complete every module. It is a sequence of revenue-gated capability releases. Each material build item must be linked to a contracted obligation, a reusable product hypothesis, and an acceptance owner.

Phase	Timing	Principal actions and exit evidence	Gate
0. Authoriza- tion	Days 0–3	Approve livestock scope, owners, claims register, currency, source hierarchy, and baseline-reconciliation plan.	Signed charte
1. Executable release	Days 1–10	Reproduce event and source paths; validate self-service guides; tag release; publish synthetic fixtures and limitations.	Clean-machin pass
2. Private acti- vation	Days 8–20	Run representative events/manifests; document gaps; identify obligations; convert reusable gaps into scopes.	Qualified map pings
3. First paid conversion	Days 15–30	Execute approved proposal process; contract one bounded sprint; link deliverables to release and acceptance tests.	Contracted re enue
4. Production and external review	Days 31–90	Harden one source; complete one profile; operate failures; as-semble and review one artifact; prove second-customer reuse.	External/reus evidence
5. Recurring productization	Months 4–6	Launch managed operations for qualified flow; establish ser-vice reporting, security controls, observability, profile/version operations, and artifact regeneration.	Recurring con tract

6. Livestock reuse layer	Months 7–12	Add reusable event/source mappings, rights/allocation controls where demanded, artifact profiles, migration tests, and integration templates.	Declining be-spoke effort
7. Expansion decision	Month 12	Evaluate New Zealand and adjacent agricultural domains only against recurring revenue, reuse, external citations, and capacity.	Executive approval

21.1 Month 4–6 product priorities

Assuming Day-90 gates are met, the next priority is not feature breadth. It is operational reliability for the first recurring flow:

- customer-specific security and access controls;
- monitoring, incident response, backup, recovery, and change management;
- source schema and mapping version operations;
- method/profile update and migration workflow;
- repeatable review queue and artifact regeneration;
- managed evidence service reports;
- customer and external-review support boundary; and
- gross-margin and support-burden measurement.

21.2 Month 7–12 product priorities

The product should expand horizontally only where reuse is demonstrated:

- common livestock identity and authority bindings;
- intervention, delivery, dosing, measurement, calibration, batch, facility, and model event families supported by paid mappings;
- reusable source-manifest profiles for trials, labs, assays, and regulatory documents;
- common evidence-gap taxonomies;
- rights, co-funding, allocation, claim-cap, reversal, retirement, and duplicate-claim controls where required;
- repeatable buyer, verifier, sponsor, and reporting artifact profiles;
- versioned Australian profiles and historical replay; and
- evidence-to-release-to-revenue attribution.

21.3 Explicitly deferred domains

Soil carbon, broad remote-sensing validation, autonomous machinery, fermentation operations, hydrogen systems, and other domain families should not dictate the initial product. They may enter the roadmap only when:

1. a named customer and external obligation exist;
2. the required vocabulary and profile can be designed without contaminating the livestock core;
3. the work produces reusable assets; and
4. the team has enough capacity to maintain current production and open/publication commitments.

Source basis: S1, S4, S5, S8..

22 Scientific Publication and Institutional Strategy

22.1 Objective

AgEvidence should establish an open, peer-reviewed reference architecture that Australian researchers, developers, verifiers, buyers, and public institutions can inspect, test, cite, and implement. Publication is a force multiplier for product adoption, not a substitute for customer implementation.

The safer external framing is a **reference-architecture strategy**, **institutional convergence strategy**, **open conformance benchmark**, or **standards-adjacent publication program**. The internal phrase “standards trap” describes a competitive effect, not a formal Australian Standard, government approval, exclusionary licensing, collusion, or restriction on lawful competition.

22.2 Publication themes

A coordinated portfolio may cover:

- country-neutral agricultural evidence envelopes;
- provenance and receipt architecture;
- separation of integrity, scientific, method, review, verification, artifact, reliance, allocation, and payment states;
- evidence portability across methodologies and institutions;
- Australian livestock production adapters;
- methane-technology-neutral evidence requirements;
- producer authority, co-funding, claim allocation, caps, reversals, and duplicate-claim controls;
- controlled source-manifest patterns for trials, labs, and biological inputs;
- reproducibility, offline reconstruction, and local verification; and
- institutional-reliance benchmarks and artifact profiles.

22.3 Strategic effects

The publication program should establish:

- common vocabulary and stable interface expectations;
- evidence-quality and gap taxonomies;
- portability and migration tests;
- reviewer and verifier expectations;
- procurement and assurance criteria;
- reusable synthetic or permissioned research fixtures;
- independent citations and implementations; and
- a lower-friction path for upstream systems to emit AgEvidence-compatible evidence.

22.4 Publication and product synchronization

The product repository, technical preprint, scientific papers, customer artifacts, and Project 4030 or other institutional outputs must not invent competing vocabularies. A shared architecture governance process should control:

- term definitions and graph boundaries;
- event and source-record semantics;

- trust-state definitions;
- country and method adapter boundaries;
- claims and prohibited inferences;
- conformance fixtures and verifier behavior;
- versioning and supersession; and
- publication-to-release traceability.

22.5 Scientific independence

Commercial leadership may set product priorities and fund implementation, but it must not predetermine scientific findings. Required controls include:

- written author and affiliation consent;
- transparent conflicts and funding statements;
- independent scientific review and documented limitations;
- reproducible methods and accessible conformance assets where rights permit;
- separation between scientific governance and commercial quote approval;
- technology-neutral evidence requirements; and
- no implication that publication establishes regulatory approval, efficacy, claim ownership, or verifier acceptance.

22.6 Twenty-four-month sequence

Timing	Institutional asset	Required outcome
Months 0–3	Technical preprint and frozen release	Explain the two entry paths, append-only projection, review, artifact, local verification, and limitations using synthetic live-stock evidence.
Months 3–6	Schema and conformance alpha	Publish core vocabulary, schemas, fixtures, verifier behavior, and prohibited-inference rules.
Months 4–12	Scientific manuscripts and reference implementations	Submit coordinated papers and independent implementations; keep customer adoption and authority claims evidence-based.
Months 6–12	Reviewer and buyer artifact pilots	Obtain structured feedback on artifact usability and evidence gaps; record bounded decisions.
Months 12–18	Benchmark and procurement assets	Release permissioned benchmarks, profile tests, and evidence-expectation templates where rights permit.
Months 18–24	Institutional adoption review	Measure citations, integrations, reviewer use, procurement references, and commercial capture; decide expansion.

Source basis: *S4*, *S6*, *S8*.

23 Success Metrics and Management Scorecard

The scorecard must separate activity, technical success, commercial status, external reliance, and strategic adoption. It must not allow sandbox activity or public attention to substitute for customer and institutional proof.

23.1 Commercial metrics

Track separately:

- sandbox quoted value;

- formal proposed value;
- contracted value;
- invoiced value;
- cash collected;
- recognized revenue;
- annual recurring revenue;
- gross margin by engagement and by managed service;
- expansion revenue; and
- revenue attributable to reusable releases, mappings, and artifact profiles.

23.2 Developer and activation metrics

- clean-machine setup completion and time;
- guide completion without private explanation;
- first valid event or manifest;
- schema-validation success;
- time to first evidence projection;
- number of private and active integrations;
- representative evidence mappings received;
- reusable mapping rate;
- support requests and support hours per activation; and
- conversion from technical action to qualified paid scope.

23.3 Product and operations metrics

- accepted versus rejected events;
- duplicate and idempotency behavior;
- projection success and exception rate;
- outbox and webhook delivery latency and failure;
- replay success and dead-letter rate;
- evidence-lineage completeness;
- artifact assembly and local verification success;
- model candidate-to-human-review disposition;
- method/profile coverage and migration-test success;
- security, incident, recovery, and access-review outcomes; and
- bespoke code and analyst reconstruction time per engagement.

23.4 Reliance metrics

- named external reviews initiated;

- artifacts independently reconstructed or inspected;
- acceptance, rejection, qualification, and request-for-more-evidence decisions recorded;
- time from source event or manifest to artifact;
- evidence gaps resolved before formal review;
- percentage of artifacts requiring manual reconstruction;
- number and type of external institutions using the package; and
- recurrence of the same artifact profile across customers.

23.5 Strategic and publication metrics

- external schema implementations;
- conformance-test use;
- independent citations and peer-review milestones;
- research and university collaborations with documented rights;
- reviewer use of the reference vocabulary;
- procurement or assurance documents referencing compatible evidence expectations;
- number of upstream platforms emitting compatible evidence; and
- ratio of open adoption to paid implementation and managed-operations conversion.

23.6 Thirty-day management targets

Metric	Target	Interpretation
Reproducible tagged campaign release	1	Technical launch integrity.
Private target accounts activated	5–8	Initial cohort coverage, not pipeline value.
External technical actions completed	3+	Real activation signal.
Self-service paths completed without private explanation	3+	Documentation-product fit.
Critical guide blockers at public launch	0	External activation integrity.
Real evidence mappings received	2+	Product-problem fit signal.
Formal proposals issued	1–3	Commercial qualification.
Paid engagement	1	Revenue thesis validated.
Customer-specific branches	0	Thin-waist architecture discipline.
Unsupported production or authority claims	0	Claims discipline.
Preprint and release package	1	Institutional distribution engine started.

23.7 Twelve-month directional outcomes

The source materials suggest a desired end state of one production livestock event family, one hardened Australian source integration, one released receipt and verifier path, two to four managed evidence customers, multiple bounded reliance decisions, a sponsor-ready portfolio pattern, declining support effort per project, and revenue attributable to reusable capabilities. These are directional ambitions to be recalibrated after the Day-30 and Day-90 gates, not forecasts.

Source basis: S1, S4, S8..

24 Governance, Ownership, and Decision Rights

24.1 Core roles

The program requires named owners for:

- executive sponsorship and final continue/revise/stop authority;
- campaign execution and account activation;
- technical release and architecture;
- integrations and Event Inbox operations;
- evidence architecture and graph governance;
- scientific and measurement review;
- model-service and reproducibility engineering;
- country/method profile ownership;
- publication and authorship governance;
- legal, claims, contracts, and external-authority review;
- data stewardship, rights, privacy, and retention;
- security and production-readiness review;
- commercial qualification, pricing, proposal, and revenue-state control;
- external-safe documentation and developer experience; and
- managed evidence operations.

The source campaign designates AgEvidence Executive Sponsor as executive sponsor, with authority over scope boundaries, target authorization, escalation, commercial/open boundary, and continue/revise/stop decisions.

24.2 Proposed RACI

Deliverable / decision	Exec. sponsor	Campaign	Tech	Integration	Science	Legal/data	Commercial	Publication
Launch charter and scope	A	R	C	C	C	C	C	I
Canonical release and capability register	I	C	A/R	R	C	C	I	C
Self-service guide validation	I	A/R	R	R	C	C	C	C
Claims and prohibited-inference matrix	A	R	C	C	R	R	C	C
Private account activation	A	R	C	R	C	C	R	I
Architecture Sprint proposal	A	C	C	R	C	C	A/R	I
Customer data and rights approval	I	C	C	C	C	A/R	C	I
Method/profile approval	I	C	C	C	A/R	C	I	C
Reliance artifact release	A	C	R	R	C	R	C	I
External review coordination	A	R	C	C	R	C	C	I
Preprint and publication	C	C	C	C	R	C	I	A/R
Day-30 and Day-90 gate	A/R	C	C	C	C	C	C	I

24.3 Decision forums

24.3.1 Weekly commercialization review

Review technical activations, customer evidence gaps, named external obligations, proposals, commercial state, reusable capabilities, delivery risk, support burden, and claims issues. The meeting should be asynchronous-first, with written updates and a short blocker session.

24.3.2 Architecture review

Approve new event types, source-manifest profiles, canonical vocabulary, graph-model changes, adapter boundaries, mapping classes, release reuse, and open-versus-proprietary placement. New event families require evidence that the existing contract and source-record path cannot represent a material state transition.

24.3.3 Scientific and claims review

Approve scientific statements, measurement assumptions, method interpretations, publication language, limitations, external authority references, customer-facing evidence claims, and prohibited inferences.

24.3.4 Data, rights, and security review

Approve source authority, permissions, retention, access, de-identification, publication rights, security posture, production use, and incident response.

24.3.5 Commercial approval

Approve qualification, price logic, scope, acceptance criteria, payment schedule, production-hardening assumptions, external-review boundary, and revenue-state treatment.

24.4 Issue taxonomy

Issue class	Examples	Required disposition
Launch blocker	Clean setup fails; canonical guide broken; example not reproducible.	Resolve before private or public launch.
Disclosed limitation	Fixture model only; partial key management; sandbox billing.	Publish clearly; may not block controlled evaluation.
Account mapping gap	Customer field or event has no current representation.	Evaluate external obligation, reuse, and paid scope.
Claims risk	Copy implies approval, verification, compliance, reliance, or customer adoption.	Remove or qualify before release.
Rights/data risk	Evidence cannot be stored, published, or referenced safely.	Use controlled path, de-identify, obtain rights, or exclude.
Architecture violation	Country fork; crypto in application layer; model treated as authority; source-platform replacement.	Reject or escalate to technical governance.
Commercial ambiguity	Sandbox confused with contract/payment/revenue; deliverables unclear.	Correct ledger and commercial documents immediately.
Operational failure	Queue, webhook, replay, dead-letter, projection, or artifact failure.	Preserve original evidence; diagnose through run-book; record incident and fix.

Source basis: S1, S4, S6, S8..

25 Risks and Controls

25.1 Overclaiming readiness

Failure mode

An implemented scaffold is presented as production-ready, secure, scientifically valid, registry-accepted, verified, or institutionally relied upon.

Primary controls: maintain the implemented/partial/absent capability register; require tagged release and test evidence; use explicit limitations; commission independent review; prohibit unsupported customer, production, authority, and revenue claims.

25.2 Validity laundering

Failure mode

Cryptographic integrity is represented as scientific validity, method compliance, verification, institutional acceptance, claim ownership, or payment legitimacy.

Primary controls: enforce separate state families in data models, APIs, interfaces, artifact language, dashboards, contracts, and metrics; prohibit ambiguous global “valid” labels.

25.3 Building before demand

Failure mode

Engineering expands into broad platform functionality before willingness to pay and external use are demonstrated.

Primary controls: require a named customer, source, obligation, contract, acceptance test, and reuse case before material hardening; use Day-30 and Day-90 stop authority.

25.4 Consultancy trap

Failure mode

Every engagement becomes bespoke professional services with no declining marginal effort.

Primary controls: require each project to produce reusable schemas, fixtures, handlers, profiles, gap taxonomies, artifact templates, runbooks, or operating capabilities; track reuse and gross margin.

25.5 Open-source value leakage

Failure mode

Commercially valuable customer-specific work is released without a capture mechanism, or the open layer is too weak to drive adoption.

Primary controls: open common grammar, schemas, reference code, tests, and selected fixtures; retain paid mappings, controlled data, production operations, customer profiles, service guarantees, and managed reliance; approve licenses and trademark rules.

25.6 Institutional or scientific capture

Failure mode

Research appears to predetermine results in AgEvidence’s favor or convert commercial priorities into scientific authority.

Primary controls: independent review, author and affiliation consent, transparent conflicts, reproducible methods, documented limitations, technology neutrality, and separate scientific governance.

25.7 Customer data and rights risk

Failure mode

AgEvidence stores, publishes, reuses, or discloses evidence without clear authority, producer consent, or contractual rights.

Primary controls: rights graph, controlled source records, explicit permissions, data classification, de-identification, retention and deletion rules, publication-rights approval, and data-steward signoff.

25.8 Target-account assumption risk

Failure mode

Public funding, staffing, deployment, or strategic announcements are mistaken for remaining budget, decision authority, or willingness to buy.

Primary controls: dated source verification; direct target validation; confirm technical sponsor, economic buyer, budget, contracting authority, external deadline, alternative cost, data rights, and acceptance tests before quoting.

25.9 Policy and adapter drift

Failure mode

Country or method rules are hard-coded into the evidence core, silently updated, or treated as universal authority.

Primary controls: declarative versioned profiles, effective dates, source basis, migration and historical replay, independent approval, and explicit externally dependent rules.

25.10 External reliance concentration

Failure mode

One buyer, verifier, registry, or sponsor profile becomes a de facto product fork or is represented as universal acceptance.

Primary controls: institution-specific bounded profiles, no evidence-history rewrite, portability tests, and reliance events that preserve decision scope and limitations.

25.11 Operational and security failure

Failure mode

Production event loss, duplication, unauthorized access, webhook compromise, key failure, or unrecoverable artifact inconsistency undermines trust.

Primary controls: threat modeling, key management, access control, monitoring, idempotency, replay, dead-letter controls, backup/recovery, incident response, release provenance, and customer-specific security acceptance.

Source basis: S1, S2, S4, S5, S6, S8..

26 Stop Conditions

The program should stop or materially revise when any of the following persists beyond the relevant gate:

- ▲ no target completes a meaningful technical action;
- ▲ no representative event or controlled source record can be mapped without a platform replacement;

- ▲ no named external reliance obligation, reviewer, buyer, or decision emerges;
- ▲ no customer will pay for the initial bounded engagement;
- ▲ the first high-value customer requires a separate product, country-specific cryptographic fork, or broad domain build;
- ▲ security, privacy, data-rights, or contractual requirements cannot be met responsibly;
- ▲ external reviewers reject the architecture for fundamental reconstruction, provenance, authority, or usability reasons;
- ▲ the implementation produces no reusable schema, handler, profile, artifact, gap taxonomy, or operating procedure;
- ▲ customers value only generic carbon accounting rather than the evidence and reliance layer;
- ▲ scientific, regulatory, or institutional claims cannot be maintained within approved boundaries;
- ▲ every artifact requires undocumented analyst reconstruction;
- ▲ recurring managed evidence operations are not credible after the first project; or
- ▲ gross margin and support burden do not improve with reuse.

Continuation rule

Continue only while each customer makes the common evidence plane more reusable, more externally inspectable or relied upon, and more profitable per unit of engineering and review effort.

Source basis: S1, S4, S8..

27 Investment Case and Financing Milestones

27.1 What seed capital must prove

Seed capital should not be justified by the breadth of the conceptual platform. It should be justified by a sequence of de-risking milestones that establishes product demand, technical reuse, institutional utility, and improving unit economics.

No.	Milestone	Evidence required	Venture significance
1	Paid problem validation	At least one contracted evidence-architecture or readiness engagement tied to a named external obligation.	Establishes willingness to pay.
2	Productized delivery	One real source flow processed through the shared evidence contract with documented acceptance criteria.	Establishes product reality beyond a demo.
3	External utility	One named reviewer or relying institution inspects a bounded artifact and records structured feedback or a decision.	Establishes institutional usefulness.
4	Reuse	A second customer reuses a material share of schemas, handlers, projections, profiles, or artifact logic.	Distinguishes infrastructure from bespoke consulting.
5	Recurring revenue	At least one customer converts to managed evidence operations or an equivalent annual service.	Establishes recurring business-model potential.
6	Margin progression	Support effort, bespoke code, and review burden decline per repeated implementation.	Establishes scalable unit economics.
7	Distribution moat	External schema use, conformance testing, citations, or institutional vocabulary adoption begins.	Establishes category and ecosystem leverage.

27.2 Use-of-funds logic

The source materials do not specify a financing amount or detailed operating budget. A VC-ready financing plan should therefore be completed during fundraising diligence. The recommended allocation logic is:

- product and integration engineering for production hardening of the Event Inbox, source-manifest path, evidence projection, policy adapters, and artifact workflows;
- security, privacy, data-rights, observability, and reproducibility controls required for real customer evidence;
- customer implementation capacity for one or two tightly bounded livestock integrations;
- scientific, assurance, legal, and claims review to preserve the separation of integrity, method compatibility, verification, and reliance;
- developer documentation, reference fixtures, open conformance assets, and publication work that create ecosystem distribution; and
- sufficient operating runway to reach the Day-90 and Month-12 revenue, reuse, and recurring-service gates without broadening the scope prematurely.

27.3 Financing conditions and investor diligence

A financing decision should be conditioned on:

- verification of the canonical repository baseline, ancestry, tests, and clean-machine reproducibility;
- a current capability register distinguishing implemented, scaffolded, and externally dependent functions;
- direct validation of Wave 1 target-account assumptions and customer budgets;
- a documented founder and operating-team plan with named technical, commercial, science, legal/data, and security ownership;
- a financing model covering runway, hiring, gross margin, implementation capacity, and follow-on capital requirements;
- approved data-rights and security controls for real evidence;
- bounded contracts with acceptance tests and no guarantee of external reliance; and
- preservation of explicit stop conditions if paid demand, reuse, institutional utility, or recurring economics fail to emerge.

28 Source Reconciliation Required Before External Finalization

The source corpus contains several conflicts and scope issues that must be resolved before this document becomes investor-final, customer-ready, or public.

28.1 Repository baselines

The materials reference multiple commits, including `f4ec679c...`, `7d315b6...`, `1b9ff053...`, and `f71d04d...`. The most recent Event Inbox decision memorandum describes `f71d04d16c747374c215734e8ce191824dd63592` as the minimum viable commercial spine, while the canonical campaign describes `1b9ff053b51b0dda2ed14a19e0f491e69629e9e4` as the orientation and self-service layer built around the underlying Developer OS.

Proposed interim reconciliation:

- treat `f71d04d...` as the working Event Inbox commercial baseline described by the C-Suite decision memorandum;
- treat `1b9ff053...` as the campaign/documentation distribution layer;
- retain earlier commits as historical implementation milestones only; and

- require the technical owner to verify repository ancestry, current head, tests, and release contents before external use.

This proposal does not claim that the interim hierarchy is technically authoritative until that verification is complete.

28.2 Campaign versions

The canonical 30-day campaign must explicitly supersede earlier campaign variants. External and internal guide surfaces should be named once and linked from one authoritative index.

28.3 Pricing frameworks

The standardized campaign offers are entry anchors. The commercial diligence models provide account-specific bounded ranges based on project value, alternative cost, budget, complexity, and risk. They are complementary, not equivalent price lists.

28.4 Currency

This proposal uses U.S. dollars as the primary memorandum currency because a U.S.-dollar commercial diligence edition was supplied. Australian-dollar account models may appear in internal diligence annexes with a dated exchange-rate basis. No mixed-currency total should be presented without an explicit conversion date and source.

28.5 Generic CPG summary

The supplied generic consumer-packaged-goods executive summary analyzes broad supermarket and consumer-product entry. It is outside the AgEvidence evidence-infrastructure thesis and is excluded unless AgEvidence separately proposes a consumer product.

28.6 Named-account assumptions

Company architecture, funding, staffing, deployment, decision authority, external obligations, and willingness-to-pay hypotheses require direct validation before external assertion. Account dossiers must distinguish verified fact, dated public claim, reasonable inference, internal hypothesis, proposed activation, and prohibited claim.

28.7 External legal and regulatory facts

Any statement concerning Australian climate disclosure, assurance, APVMA, registry, corporate, privacy, data-localization, or other legal and regulatory obligations must be refreshed from current primary sources and approved by counsel before external use. This proposal retains the source-derived strategic framing but does not independently verify those external facts.

Non-negotiable guardrail

The resulting memorandum is a milestone-gated venture investment case. It is not proof that the market, customers, scientific claims, regulatory acceptance, verifier acceptance, institutional reliance, recurring operations, or revenue already exist.

Source basis: S1–S9; reconciliation synthesized from the supplied outline and source corpus..

A Ninety-Day Acceptance Criteria

A.1 Release acceptance

Criterion	Evidence required	Approver
Canonical release identified	Exact SHA/tag, repository URL, ancestry note, release manifest, schema/event versions.	Technical owner
Clean-machine reproducibility	Installation log, automated tests, demo script, expected outputs, failure examples, local verification.	Technical + operations
Capability register verified	Implemented/partial/absent matrix with evidence and owners.	Architecture review
Guide pack validated	Each guide executed against release; blockers resolved or disclosed.	Campaign lead
Claims matrix approved	Permitted claims, required qualifiers, prohibited inferences, source status.	Legal/science/exec.

A.2 First customer sprint acceptance

Criterion	Evidence required	Approver
Representative source mapped	Real or customer-approved controlled event/source manifest; source authority and rights documented.	Customer technical owner
Evidence lineage complete	Source-to-projection-to-review-to-artifact links reproducible.	AgEvidence evidence owner
Gap report delivered	Findings, severity, owner, remediation, external-obligation relevance, and limitations.	Customer + AgEvidence
Review decisions recorded	Reviewer identity, authority, rationale, evidence version, and disposition.	Scientific/evidence owner
Profile applied	Versioned method/country/institution profile and compatibility result with external dependencies.	Profile owner
Artifact verified	Manifest complete; local verification passes independently; limitations visible.	External reviewer / QA
Commercial acceptance	Named deliverables accepted under contract; invoice/payment state recorded separately.	Customer economic buyer
Reusable asset produced	Schema, handler, fixture, profile, template, or runbook approved for common plane.	Architecture review

A.3 Day-90 continuation acceptance

- ✓ at least one contracted and delivered paid engagement;
- ✓ one named external reviewer or institution has inspected the artifact and recorded a bounded outcome;
- ✓ one second-customer or second-use-case reuse test has been completed;
- ✓ a measured managed-operations need and commercial path exist;
- ✓ security, data-rights, scientific, legal, and claims findings are acceptable or remediable;
- ✓ the product can continue without creating a customer-specific fork or replacing the customer's operating platform; and
- ✓ the board and lead investors approve the 12-month plan and budget based on evidence rather than pipeline aspiration.

B Product Requirements Traceability Matrix

Requirement	Customer job	Primary module	Acceptance evidence	Phase
Preserve upstream system authority	Connect without replacement	Event Inbox / Source Manager	Source and authority register; no system migration.	0–1
Append-only evidence history	Provenance and integrity	Inbox / Projection	Replay and history tests; no destructive update.	1
Controlled source references	Connect document evidence	Source Manifest Manager	Hash, location, authority, access, version, limitations.	1
Identity continuity	Provenance and integrity	Projection Engine	Entity correlation and conflict report.	1–2
Evidence-gap detection	Pre-review readiness	Gap Analyzer	Fixture-backed gap report and human disposition.	1–2
Human review non-authority of model	Pre-review readiness	Review Ledger	Append-only decisions; model candidates remain review-required.	1
Versioned policy evaluation	Multi-policy portability	Adapter Layer	Side-by-side profile results without evidence rewrite.	2–4
Rights and allocation separation	Multi-party claims	Rights Engine	Rights/allocation events and duplicate-claim controls.	4–6
Portable bounded artifact	Institutional reliance	Artifact Builder	Manifest, lineage, limitations, local verification.	2–4
Separate external reliance event	Institutional reliance	Reliance Registry	Named institution, purpose, decision, limits, evidence version.	3–4
Recurring operations	Operate evidence flow	Managed Evidence Ops	SLOs, runbooks, service reports, change and incident records.	5
Release-to-revenue lineage	Commercial control	Cross-module	SHA, source, engagement, artifact, decision, cash and revenue states.	2–5

C Wave 1 Activation Plays

C.1 DIT AgTech

Hypothesis: event-first integration for intervention delivery, dosing, water-flow, device, or exception evidence.

First action: map one representative operational event through the signed inbox and inspect source identity, device/software version, intervention/batch identity, units, idempotency, exception, and projection lineage.

Likely gaps: production signing and key lifecycle; event ontology; device/source lineage; calibration; identity continuity; exception semantics; external buyer or verifier profile.

Paid wedge: Evidence Architecture Sprint leading to a livestock event bridge and production-hardening scope.

Validation required: actual architecture, deployment, source owner, external obligation, budget, contracting authority, data rights, and willingness to pay.

C.2 MEQ Solutions

Hypothesis: event-first processor and measurement evidence linking device/model versions, facility context, animal or carcass identity, verification status, and buyer decision.

First action: map one measurement or model-completed event and assemble a minimal buyer/auditor evidence trace.

Likely gaps: device and model versioning; facility identity; calibration; identity reconciliation; buyer state; verifier findings; settlement and reliance linkage.

Paid wedge: Enterprise Reliance Artifact architecture, potentially followed by a production processor evidence bridge.

C.3 Agscent

Hypothesis: source-record-first evidence for measurement reports, device calibration, cohort linkage, and model interpretation.

First action: submit one controlled measurement report and model reference, run candidate/gap processing, and record human review.

Likely gaps: source provenance; calibration and device identity; cohort and intervention linkage; model-domain and uncertainty statements; external reviewer requirements.

Paid wedge: Verification Readiness Cycle.

C.4 Cibo Labs

Hypothesis: controlled spatial source and model-run evidence requiring explicit temporal bounds, raster/source commitments, model version, and limitations.

First action: add one spatial output reference and inspect the source-to-model-to-artifact chain.

Likely gaps: geospatial commitments, versioning, temporal alignment, source rights, model-domain boundaries, and method-specific interpretation.

Paid wedge: remote-sensing evidence architecture only when reuse and obligation justify the additional domain work.

C.5 Rumin8

Hypothesis: source-record-first trial, chemistry, dosing, product, batch, efficacy, safety, regulatory, and commercial-deployment evidence.

First action: submit one controlled trial, product, or regulatory reference and map product/intervention identity into the evidence graph.

Likely gaps: cross-trial provenance; batch and dose identity; cohort and delivery context; safety and efficacy boundaries; regulatory-document authority; method-profile requirements.

Paid wedge: Protocol Evidence Implementation.

C.6 Sea Forest

Hypothesis: source-record-first feed-additive evidence connecting product authorization, supply and batch, feeding, herd/cohort, measurement, rights, and Scope 3 claim allocation.

First action: submit one feed-additive evidence reference and map trial-to-field evidence lineage.

Likely gaps: supply/batch continuity; administration and cohort evidence; claims rights; co-funding and allocation; buyer artifact requirements; duplicate-claim controls.

Paid wedge: Scope 3 reliance architecture and artifact.

Non-negotiable guardrail

All Wave 1 descriptions are internal activation hypotheses. They must not be represented as customer architecture, budget, intent, deployment, efficacy, or commitment without direct validation and dated sources.

D Source Basis and Claims Register

D.1 Source register

ID	Source	Primary contribution to this proposal
S1	<i>AgEvidence Australia Event Inbox C-Suite Decision for Kendra</i>	Working Event Inbox commercial spine; bridge-first 90-day decision; capability boundaries; stop conditions; ownership and external-review sequence.
S2	<i>AgEvidence ANZ Institutional Commercial Diligence - U.S. Dollar Edition</i>	Bounded value pricing; revenue-state definitions; Wave 1 economics; account diligence; U.S.-dollar planning basis.
S3	<i>AgEvidence ANZ Institutional Commercial Diligence - Australian Dollar Edition</i>	Account-specific Australian-dollar pricing ranges and commercial diligence context.
S4	<i>The Zero Sales Team Campaign - Canonical 30-Day Developer-Led Sprint</i>	Developer-led activation, self-service distribution, repository guides, product ladder, metrics, RACI, claims discipline, and Day-30 gate.
S5	<i>Canonical Australian Target Account Dossiers</i>	Wave 1 account sequence, technical activation hypotheses, evidence mapping classes, and target-validation requirements.
S6	<i>Institutionalizing the Competitive Standard</i>	Open evidence commons, institutional convergence, five-graph model, three-tier architecture, publication strategy, governance, and claims boundaries.
S7	<i>Australian AgEvidence Product Proposal and Commercial Strategy Research</i>	Market and product synthesis, dual ingestion paths, trust-state model, graph decomposition, account scoring, and commercial architecture.
S8	<i>AgEvidence Australia Comprehensive Product Proposal Outline</i>	Requested proposal structure, decisions, modules, packaging, roadmap, governance, risks, reconciliation, and drafting constraints.
S9	<i>PP Executive Summary</i>	Reviewed and excluded from the core thesis because it focuses on consumer packaged goods and supermarket entry.

D.2 Claims classification

Every material statement used externally should be classified as one of the following:

Class	Definition and use rule
-------	-------------------------

Verified repository fact	Confirmed against the selected tagged release, tests, and file path. May be stated precisely with release reference.
Verified external fact	Confirmed against a current primary source and approved for external use. Must include date where material.
Customer-validated fact	Confirmed directly by an authorized target representative and documented. Use within agreed confidentiality.
Reasonable inference	Derived from verified facts but not directly confirmed. Must be labeled as inference and not presented as commitment.
Internal hypothesis	Proposed account fit, budget, architecture, product wedge, or outcome. Internal use only until validated.
Management proposal	A recommended future decision, process, price, role, or roadmap. Must not be described as implemented or approved.
Prohibited claim	Unsupported statement of approval, verification, customer adoption, efficacy, authority, revenue, production readiness, or universal acceptance. Must not be used.

D.3 Core prohibited claims

Unless separately evidenced and approved, the proposal and campaign may not claim:

- production certification or enterprise security approval;
- APVMA, registry, regulator, government, buyer, or verifier approval;
- verified methane reduction or scientific efficacy of a target technology;
- customer adoption, pilot participation, willingness to pay, contract, payment, or revenue;
- universal method compatibility or claim ownership;
- that a country adapter has government authority;
- that publication creates a formal Australian Standard; or
- that one institutional decision binds other institutions.

E Glossary

Term	Definition
Evidence contract	The stable canonical envelope and semantics through which source systems commit evidence and downstream processes interpret it.
Evidence Event Inbox	Append-only intake path for signed operational events from existing systems.
Controlled source manifest	Hash, location/reference, authority, version, and metadata describing a source file retained in controlled infrastructure.
Evidence projection	Canonical representation derived from an accepted source event or manifest while preserving lineage.
Evidence gap	Missing, inconsistent, unsupported, uncalibrated, unauthorized, or otherwise insufficient information relevant to a named obligation.
Human Review Ledger	Append-only record of authorized reviewer decisions, rationale, evidence version, and limitations.
Policy/method adapter	Versioned profile that evaluates preserved evidence under country, method, buyer, assurance, registry, or contractual requirements.
Reliance artifact	Portable package assembled for a named institution and decision, with manifest, lineage, review, profile, limitations, and verification instructions.
Reliance event	Separate record of an external institution's acceptance, rejection, qualification, or request for more evidence.

Thin waist	Small stable contract connecting many upstream source systems with many downstream methods and institutions without requiring one monolithic stack.
Validity laundering	Improperly treating one valid state, such as signature integrity, as proof of science, compliance, verification, reliance, ownership, or payment legitimacy.
Open evidence commons	Public vocabulary, schemas, fixtures, tests, reference tooling, and research assets that lower compatibility and evaluation friction.
Managed Evidence Plane	Recurring service for ingestion, exception handling, profile updates, review operations, artifact regeneration, support, and reliance records.
Bounded engagement	Paid scope tied to a named source, evidence gap, external obligation, deliverable, acceptance test, and reusable capability.

F Final Recommendation

AgEvidence should authorize a narrow, disciplined Australia-first product commercialization test around the working evidence bridge already described in the repository materials. The goal is not to declare a complete platform. The goal is to prove that the least disruptive integration path can also become the most defensible evidence path.

The correct sequence is:

Working bridge → paid mapping → one hardened source → bounded external artifact → reusable second customer → managed evidence revenue

AgEvidence should continue only when real customers, real evidence, real external obligations, reusable product increments, and bounded institutional decisions justify the next stage. It should stop when the market requires a replacement platform, a country-specific core fork, unsupported authority claims, or bespoke work that does not improve the common evidence plane.

Approve the bridge-first 30-day proof and 90-day commercial test. Preserve the source systems. Separate every trust state. Sell bounded evidence and reliance work. Build only what contracted evidence and external review prove should become reusable product.

Document status: Confidential venture investment memorandum. Prepared from the supplied source materials for prospective investors and the AgEvidence founding team. All external claims require dated verification, rights review, and approval by the relevant technical, scientific, legal, data, security, commercial, and executive owners.