

PRIVATE INSTITUTIONAL PREVIEW • v13.0.0-GL1

GoalOS Global Launch Ω

The Verified Operating System for Building Anything, Anywhere.

Tell GoalOS what you want to build.

It searches the world for the strongest defensible institutional architecture; coordinates entities, markets, talent, capital, public support, capability, procurement and infrastructure; launches bounded AGI Jobs; blocks incomplete transactions; proves what happened; and keeps the institution current.

SEARCH BROADLY

Generate diverse whole-institution architectures and non-obvious positive-upside paths.

EXECUTE EXACTLY

Constrain data, tools, budget, authority and external action through fail-closed gates.

ADMIT SELECTIVELY

Only rights-cleared, evidence-supported and independently challenged capability enters Chronicle.

COMPOUND VERIFIED VALUE

Reinvest realized contribution into productive capacity only after the complete denominator passes.

Ω Apex Opportunity Search

Broad search. Exact authority. Positive-upside frontier.

The v13 flagship extends the proof-governed architecture with a bounded opportunity-search council. It allocates additional search and simulation to the consequential uncertainty, not to uncontrolled external action.

1

CONSTITUTE

Freeze objective, sponsor, authority, budget, deadline, success and prohibitions.

2

GENERATE

Create concentrated, hybrid, market-first, R&D-first, infrastructure-first and distributed architectures.

3

SIMULATE

Evaluate thousands of worlds across funding delay, tax friction, market access, compute, execution and concentration shocks.

4

CHALLENGE

Attack substance, permanent establishment, sanctions, data, grants, rights, conflicts, clawbacks and unwind cost.

5

PROVE

Expose Proof Debt, source expiry, professional dependencies, transaction receipts and independent-verifier requirements.

6

ADMIT

Publish a non-dominated frontier and select the highest verified-upside scenario that remains lawful and reversible.

7

EXECUTE

Launch bounded AGI Jobs and authorized provider work only within exact scope.

8

LEARN

Admit narrow capability to Chronicle only after realized outcomes and a fresh successor comparison.

APEX SEARCH LAW

Candidate intelligence may search expansively. Institutional action remains fail-closed. No model-generated result becomes legal, tax, regulatory, fiduciary, financial, employment, procurement, public or infrastructure authority without current facts, attributable evidence and responsible approval.

THE WORLD IS THE LAUNCH STACK. THE EVIDENCE DOCKET IS THE GATE. REALITY IS THE FINAL EVALUATOR.

MONTREAL.AI RESEARCH

PROOF-GOVERNED INTELLIGENCE INSTITUTIONS

GoalOS Global Launch Ω

The Verified Operating System for Building Anything, Anywhere



Tell GoalOS what you want to build.

It turns the world into your launch stack—and the objective into a verified operating institution.

Version 1.0.0 | July 2026

Preprint and institutional design paper

This paper introduces *GoalOS Global Launch Ω*, a proof-governed operating system for forming, financing, deploying, and continuously improving institutions across jurisdictions. Existing launch markets are fragmented across company formation, employer-of-record services, cap-table administration, legal and tax research, public-funding discovery, procurement, site selection, and professional advisory. GoalOS Global Launch treats these as modules of a single consequential objective: determining how a lawful entity, programme, fund, consortium, expansion, infrastructure project, or economic-development initiative should be architected and executed worldwide. The system decomposes a launch objective into eight independently optimized functions—sell, hire, perform R&D, own capability, raise capital, obtain public support, deploy infrastructure, and realize value—then generates competing institutional scenarios, assigns bounded AGI Jobs, calculates capital and public-support stacks, exposes unresolved Proof Debt, and blocks material transactions until required evidence and responsible authority exist. Completed work is recorded in Evidence Dockets; only narrow, rights-cleared, independently supported capability may enter Chronicle and influence a successor generation. We formalize the launch mission, scenario frontier, proof-gated transaction model, verified launch value function, and recursive improvement test; describe ten natural product modules; present commercialization and recurring-monitoring economics; and define an empirical proof programme. The central claim is not that a model can replace lawyers, tax advisers, regulators, fiduciaries, or public authorities. It is that global institution formation can be reorganized around explicit objectives, graph-structured work, bounded authority, attributable evidence, selective memory, and measured improvement.

Keywords: agentic AI; institution formation; jurisdiction intelligence; public funding; fiscal optimization; AGI Jobs; proof governance; graph engineering; capital allocation; recursive improvement.

Claim Boundary

This paper presents a testable institutional architecture and commercialization thesis. It does not claim that GoalOS Global Launch has achieved product-market fit, binding legal or tax authority, universal jurisdictional coverage, guaranteed public funding, autonomous regulated practice, or civilization-scale productive capacity. Any binding formation, filing, transfer, financing, employment, tax, regulatory, or infrastructure action requires current official sources, exact facts, qualified responsible professionals, and scoped human authority.

Contents

1. Introduction	4
2. The Fragmented Launch Stack	4
2.1. Point solutions optimize local tasks	4
2.2. Why an AI wrapper is insufficient	5
3. Category Definition: Global Institutional Launch Intelligence	5
3.1. Objects that may be launched	6
3.2. The eight functions	6
4. Formal Problem Statement	6
4.1. Launch mission	6
4.2. Institutional scenario	7
4.3. Verified launch value	7
4.4. Proof Debt	8
4.5. Transaction gate	8
5. System Architecture	8
5.1. GoalOS Chat Ω	10
5.2. Launch Constitution	10

5.3. Jurisdiction auction and scenario frontier	10
5.4. AGI Job graph	10
5.5. Evidence Docket and Chronicle	10
6. Natural Expansion Modules	11
7. Graph Engineering and Separation of Powers	12
7.1. Five edge types	12
7.2. Independent verification	12
7.3. Institutional planes	12
8. Capital and Public-Support Architecture	13
8.1. Funding order	13
8.2. Stacking and anti-double-dip engine	13
9. Governance, Safety, and Legal Boundary	13
9.1. No secrets before contract	14
9.2. Professional authority	14
9.3. Fail-closed execution	14
10. Recurring Operation: The Sovereign Expansion Office	14
11. Recursive Self-Improvement	14
11.1. Chronicle admission	14
11.2. Fresh successor test	15
12. Illustrative Mission	15
13. Commercialization Architecture	15
13.1. Product ladder	15
13.2. Customer value gate	16
13.3. Decacorn and megacorn scenarios	16
14. Empirical Proof Programme	16
15. Risks and Failure Modes	17
16. Research Agenda	17
17. Conclusion	17
A. Minimum Launch Mission Contract	19
B. Evidence Docket Schema	19
C. Canonical Operating Doctrine	19
References	20

1. Introduction

Launching an institution is not one task. It is a coupled system of market access, legal formation, ownership, talent, research, capability, financing, public support, procurement, infrastructure, regulatory authority, and continuing governance. Yet the software and advisory market presents these components as isolated products. Company-formation platforms can automate a specific incorporation path; employer-of-record platforms can hire workers in jurisdictions where a customer lacks an entity; cap-table platforms can administer equity and private capital; professional databases can retrieve law and guidance; grant platforms can surface programmes; and advisers can study bespoke questions. Each component is useful, but no component begins from the entire objective.

The fragmentation is visible in the maturity of existing point solutions. Stripe Atlas reports that more than 100,000 founders have used its Delaware incorporation workflow and that founders in more than 140 countries have adopted the service (Stripe, 2026). Remote advertises employer-of-record coverage in more than 90 countries, while Deel reports owned infrastructure and entities in more than 100 countries (Remote, 2026; Deel, 2025). Carta combines cap-table management, financing, valuations, compliance, contracts, and fund administration in one private-capital platform (Carta, 2026). These systems demonstrate substantial demand for individual launch functions. They also reveal the category gap: a sponsor still must determine which combination should exist, where each function belongs, what should happen first, which assumptions remain unresolved, and which transactions may proceed.

**The category is not business formation.
It is institution formation.**

GoalOS Global Launch begins with a different question:

What must be built, under whose authority, with which capital, in which jurisdictions, through which sequence of bounded missions, and with what evidence that the resulting institution works?

The paper makes six contributions:

1. It defines a *functional jurisdiction architecture* in which eight launch functions are optimized independently rather than forced into one jurisdiction.
2. It introduces a *scenario frontier* that generates concentrated, hybrid, and distributed institutional designs rather than a single opaque recommendation.
3. It formalizes a *proof-gated transaction model*: recommendations do not become authority; transactions remain blocked until required evidence and responsible approvals exist.
4. It integrates *AGI Jobs* and graph engineering into launch execution, separating independent work, true dependencies, verification edges, authority edges, and governance edges.
5. It defines a *Chronicle and successor mechanism*: completed work does not automatically become institutional learning; only admitted capability may influence a fresh generation.
6. It presents a commercialization model in which one-time launch missions convert into a recurring Sovereign Expansion Office and, eventually, protected capability and network economics.

2. The Fragmented Launch Stack

2.1. Point solutions optimize local tasks

Existing applications generally optimize a local object:

The problem is not that these products are weak. The problem is that their optimization boundaries are narrower than the sponsor's objective. In the language of transaction-cost economics, the launch sponsor bears the

Table 1: The fragmented launch market and the GoalOS reframing.

Application class	Primary object	Question left unresolved
Formation platform	Entity filing and standard documents	Should this entity exist, in this jurisdiction, at this time?
Employer of record	Employment without a local entity	Should the organization use an EOR, contractor, branch, or subsidiary?
Cap-table platform	Equity, financing, valuations, administration	Which capital should be raised, by which entity, under what control architecture?
Grant and incentive database	Programme discovery	Is the entity eligible, can the support stack, and when will cash arrive?
Professional research tool	Sources and commentary	Which conclusion applies to the exact facts, and who may authorize it?
Site-selection platform	Location comparison	How do market, energy, permits, financing, talent, tax, and reversibility interact?
Consulting engagement	Bespoke analysis	How can accepted work become reusable, current, and less founder-dependent?

coordination cost between specialized providers and must reconcile their incompatible assumptions, incentives, data, deadlines, and standards (Coase, 1937; Williamson, 1985).

2.2. Why an AI wrapper is insufficient

A large language model can summarize jurisdictions, draft incorporation checklists, identify grants, compare tax rates, and produce an attractive launch plan. That does not create a defensible operating institution. Plausible candidate output does not provide:

- current authoritative law or programme status;
- fact-specific eligibility;
- licensed or regulated professional judgment;
- transaction authority;
- evidence that filings, awards, contracts, hires, permits, and infrastructure actually occurred;
- complete cost, tax, dilution, delay, risk, and unwind economics;
- rights to reuse the resulting method;
- a mechanism for revocation when the world changes.

This distinction is consistent with the GoalOS Capability–Price Lag thesis: AI sharply reduces the cost of candidate cognitive output, but truth, authority, execution, accountability, and rights remain scarce (MONTREAL.AI Research, 2026c).

3. Category Definition: Global Institutional Launch Intelligence

GoalOS Global Launch Ω is defined as:

A continuously governed institution that transforms a consequential objective into a global functional architecture, capital and public-support stack, graph of bounded work, evidence-gated transactions, selective institutional memory, and continuing operating control.

3.1. Objects that may be launched

The system is intentionally broader than startups. It may support lawful:

- companies and international expansions;
- nonprofits and foundations;
- research programmes and consortia;
- investment funds and special-purpose vehicles;
- infrastructure and industrial projects;
- public institutions and economic-development initiatives;
- cooperative or multi-stakeholder institutions.

3.2. The eight functions

Let the set of launch functions be

$$\mathcal{F} = \{f_{\text{sell}}, f_{\text{hire}}, f_{\text{rnd}}, f_{\text{cap}}, f_{\text{capital}}, f_{\text{support}}, f_{\text{infra}}, f_{\text{value}}\}.$$

They correspond to:

Table 2: The eight independently optimized launch functions.

Function	Primary question	Typical evidence
Sell	Where should customers be reached, contracted, billed, and supported?	Demand, procurement access, tax registration, language, channel economics
Hire	Where and how should talent be engaged?	Labour cost, talent depth, worker classification, payroll, immigration, IP assignment
Perform R&D	Where should experimental development occur?	Talent, facilities, incentives, project evidence, nexus, collaboration
Own capability	Which entity should own or control reusable capability and IP?	Substance, assignments, rights, valuation, DEMPE, funding conditions
Raise capital	Which entity should raise each class of capital?	Investor access, control, securities rules, dilution, screening, currency
Obtain public support	Which instruments can lawfully support which costs?	Eligibility, cost share, stacking, cash timing, clawbacks, evidence
Deploy infrastructure	Where should compute, laboratories, facilities, or industrial assets be built?	Energy, land, permits, logistics, financing, offtake, resilience
Realize value	How should returns, distributions, royalties, exits, or reinvestment occur?	Tax, withholding, treaty, governance, liquidity, productive reinvestment

The best jurisdiction for one function need not be the best for another. The system therefore replaces the question “Which country is best?” with “Which architecture is most defensible for this precise mission?”

4. Formal Problem Statement

4.1. Launch mission

A launch mission is represented as

$$\mathcal{M} = (O, S, B, \tau, R, A, \mathcal{F}, \Pi, \mathcal{E}),$$

where:

- O is the consequential objective;
- S is the accountable sponsor and stakeholder set;
- B is the approved budget and capital availability;
- τ is the time horizon and deadlines;
- R is the risk, control, and reversibility posture;
- A is the authority and permission model;
- $\mathcal{F}$ is the set of functions to optimize;
- Π is the policy and regulatory constraint set;
- $\mathcal{E}$ is the evidence available at mission start.

4.2. Institutional scenario

A candidate institutional scenario s is

$$s = (r, \mathcal{J}, \mathcal{L}, \mathcal{K}, \mathcal{P}, \mathcal{T}, \mathcal{G}, Q),$$

where:

- r is the institutional root;
- $\mathcal{J} : \mathcal{F} \rightarrow \mathcal{C}$ assigns functions to jurisdictions in candidate set $\mathcal{C}$;
- $\mathcal{L}$ is the legal-entity and ownership graph;
- $\mathcal{K}$ is the capital sequence;
- $\mathcal{P}$ is the public-support portfolio;
- $\mathcal{T}$ is the controlled transaction set;
- $\mathcal{G}$ is the AGI Job and dependency graph;
- Q is the monitoring, rollback, and continuing-governance plan.

4.3. Verified launch value

The objective is not minimum tax or maximum gross subsidy. Define the risk-adjusted verified launch value:

$$\text{VLV}(s) = V_{\text{market}}(s) + V_{\text{capital}}(s) + V_{\text{support}}(s) + V_{\text{talent}}(s) \quad (1)$$

$$+ V_{\text{infra}}(s) + V_{\text{tax}}(s) + V_{\text{capability}}(s) \quad (2)$$

$$- C_{\text{formation}}(s) - C_{\text{operating}}(s) - C_{\text{coord}}(s) - C_{\text{compliance}}(s) \quad (3)$$

$$- C_{\text{dilution}}(s) - C_{\text{delay}}(s) - C_{\text{risk}}(s) - C_{\text{unwind}}(s). \quad (4)$$

The recommended scenario is not simply $\arg \max_s \text{VLV}(s)$. It is the highest-valued scenario among those satisfying evidence and authority constraints:

$$s^* = \arg \max_{s \in S} \text{VLV}(s) \quad \text{subject to} \quad \text{Gate}(s) = 1,$$

where $\text{Gate}(s) = 1$ only if all material proof and authority requirements have passed.

4.4. Proof Debt

For scenario s , let

$$\mathbf{d}(s) = \{d_1, d_2, \dots, d_n\}$$

be the set of unresolved claims, missing evidence, ambiguous authority, stale sources, regulatory dependencies, funding uncertainties, or transaction conditions. A scenario may be strategically attractive while still carrying material Proof Debt. GoalOS must expose this distinction rather than compress it into one confidence score.

4.5. Transaction gate

For a material transaction t , define the required receipt set:

$$\mathcal{R}(t) = \{r_{\text{purpose}}, r_{\text{legal}}, r_{\text{tax}}, r_{\text{valuation}}, r_{\text{tip}}, r_{\text{funding}}, r_{\text{board}}, r_{\text{implementation}}, r_{\text{chronicle}}\}.$$

The transaction gate is

$$\text{TG}(t) = \begin{cases} 1, & \text{if every applicable receipt is current, attributable, and approved,} \\ 0, & \text{otherwise.} \end{cases}$$

Transaction Law

No entity, IP transfer, licence, financing, option grant, intercompany agreement, regulated filing, major procurement, or infrastructure SPV should proceed merely because a model recommends it. If $\text{TG}(t) = 0$, the transaction remains blocked.

5. System Architecture




Figure 1: GoalOS Global Launch Ω : proof-governed architecture from objective to continuing institution.

5.1. GoalOS Chat Ω

The conversational front door collects only decision-changing, non-confidential facts: the objective, sponsor, ownership constraints, markets, capital, capabilities, deadlines, regulatory sensitivity, and risk posture. It should not ask the user to choose a jurisdiction before analysis; doing so would encode the answer into the input.

5.2. Launch Constitution

The Constitution freezes the mission boundary:

- accountable sponsor and decision authority;
- budget, time, and capacity limits;
- permitted information and tools;
- prohibited actions and data;
- acceptance criteria;
- publication and confidentiality rules;
- stop, escalation, and rollback conditions.

5.3. Jurisdiction auction and scenario frontier

The system scores jurisdictions separately for each function, then constructs multiple architectures. A single weighted average is inadequate because a high market-access score cannot erase a fatal funding restriction or regulatory incompatibility. The scenario frontier should therefore preserve non-dominated alternatives and show explicit trade-offs.

5.4. AGI Job graph

Current agent frameworks support tool use, handoffs, guardrails, tracing, independent contexts, and parallel agent teams (OpenAI, 2026; Anthropic, 2026). GoalOS uses these primitives inside a stricter institutional model. Each AGI Job has:

- a claim or decision it supports;
- authorized inputs and data scope;
- allowed tools and environments;
- budget and time limit;
- explicit output schema;
- acceptance tests;
- independent verifier where material;
- failure and escalation route;
- Evidence Docket return path.

5.5. Evidence Docket and Chronicle

The Evidence Docket records the complete denominator:

- sources and retrieval dates;
- facts, assumptions, and contradictions;
- models, tools, and versions;

- jobs, costs, retries, and failures;
- human and professional interventions;
- approvals, signatures, receipts, and acceptance;
- realized outcomes and delayed consequences;
- rights, exclusions, and reuse conditions.

Completed work is not automatically institutional learning. Chronicle admits only a narrow capability that is supported, rights-cleared, scoped, current, attributable, independently challenged, revocable, and tested on a fresh mission. This follows the broader GoalOS institutional chain from intelligence to work, proof, authority, protected capability, recurring value, capital, and productive capacity (MONTREAL.AI Research, 2026b).

6. Natural Expansion Modules

Table 3: Natural GoalOS Global Launch modules.

Module	Principal mission	Representative AGI Jobs and proof objects
GJFFI Ω	Assign functions across jurisdictions and identify funding, fiscal, regulatory, and infrastructure routes.	Source retrieval; jurisdiction scoring; funding eligibility; treaty and tax issue maps; scenario frontier; source-freshness controls.
Entity Formation Mission	Create only the entities justified by the accepted architecture.	Entity-type comparison; ownership map; governance instructions; filing checklist; banking/KYC dossier; evidence of formation.
Talent & Employment Mission	Choose between employment, EOR, contractor, branch, subsidiary, or partnership routes.	Labour-market analysis; cost model; immigration; payroll; benefits; classification; IP assignment; onboarding verification.
Capital Formation Mission	Sequence customer cash, credits, grants, debt, equity, strategic capital, and project finance.	Capital need; investor map; dilution; control; data room; term comparison; closing checklist; collected-cash receipts.
Public-Support Mission	Build a lawful, auditable support portfolio.	Programme discovery; eligibility; cost share; stacking; assistance interaction; cash timing; application dossier; award and clawback monitoring.
Capability & IP Mission	Determine creation, ownership, rights, invocation, transfer, and Chronicle admission.	Assignment audit; funding conditions; valuation; DEMPE; intercompany terms; rights register; capability scope and revocation.
Procurement & Vendor Mission	Procure the launch stack and infrastructure under policy.	Requirements; vendor discovery; RFI/RFP; security review; total cost; negotiation; contract routing; delivery and invoice verification.
Infrastructure Deployment Mission	Turn demand into land, energy, compute, permits, construction, and productive output.	Site auction; power economics; permits; logistics; vendor graph; project finance; milestone gates; commissioning tests.
Market-Entry Mission	Prove demand through the least irreversible viable route.	Market research; distribution; localization; indirect tax; payments; partner strategy; remote sales; entity trigger; customer evidence.

Module	Principal mission	Representative AGI Jobs and proof objects
Sovereign Office	Expansion Keep the institution current after launch.	Weekly source refresh; monthly funding committee; quarterly jurisdiction auction; annual substance and transfer-pricing review; trigger alerts; rollback.

7. Graph Engineering and Separation of Powers

7.1. Five edge types

GoalOS distinguishes five edge types:

- 1. **Data edge:** a downstream node consumes an upstream output.
- 2. **Proof edge:** a separate node tests a material claim.
- 3. **Authority edge:** the next stage cannot proceed without responsible permission.
- 4. **Resource edge:** jobs share or compete for state, budget, workspace, or capital.
- 5. **Governance edge:** a result cannot influence future missions without Chronicle admission.

This distinction prevents the common mistake of treating every “and then” as a dependency while also recognizing that some waits are constitutional rather than computational.

7.2. Independent verification

A verifier should receive the artifact, the acceptance criterion, and the authorized evidence—not merely the producing agent’s narrative. Where feasible, comparison should be blind and evaluators should be selected before execution. The architecture is designed to reduce Goodhart- and Campbell-style metric gaming by pairing optimizing loops with independent reality anchors (Goodhart, 1975; Campbell, 1979).

7.3. Institutional planes

Table 4: Separation of powers in GoalOS Global Launch.

Plane	May do	May not do alone
Intelligence	Research, model, simulate, compare, propose, design tests	Create external legal, fiscal, fiduciary, or regulatory authority
Execution	Perform approved AGI Jobs and coordinate providers	Exceed scope, budget, tools, data, or permission
Proof	Record tests, evidence, costs, failures, and outcomes	Convert producer assertion into independent truth
Authority	Approve, condition, pause, decline, revoke, or sign	Hide the basis of authority or bypass mandatory law
Memory	Admit, scope, version, expire, revoke, and compose capability	Store every output as truth or reuse customer secrets
Capital	Allocate verified contribution to capability and infrastructure	Reinvest booked revenue or speculative value as verified free cash flow

Institutional Law

GoalOS may propose. GoalOS may execute. GoalOS may generate a successor. Evidence must record. Independent review must challenge. Authority must approve. Chronicle must govern reuse. Reality must decide.

8. Capital and Public-Support Architecture**8.1. Funding order**

The recommended default sequence is:

$$\begin{aligned} &\text{Customer Cash} \rightarrow \text{Earned Refundable Credits} \rightarrow \text{Awarded Grants} \\ &\quad \rightarrow \text{Approved Debt} \rightarrow \text{Equity} \rightarrow \text{Strategic Capital} \\ &\quad \rightarrow \text{Project Finance.} \end{aligned} \tag{5}$$

The sequence is not universal; it is a disciplined prior. The system must separate:

- available capital;
- eligible capital;
- approved capital;
- committed capital;
- collected capital.

Runway Law

No unsigned investment, unawarded grant, unassessed tax credit, unapproved loan, or speculative customer contract enters the base runway.

8.2. Stacking and anti-double-dip engine

For cost object c and candidate instrument p , define net support value:

$$\text{NSV}(c, p) = G(c, p) - A(c, p) - T(c, p) - C_{\text{compliance}}(c, p) \tag{6}$$

$$- C_{\text{delay}}(c, p) - R_{\text{audit}}(c, p) - R_{\text{clawback}}(c, p), \tag{7}$$

where G is gross support, A is assistance interaction, T is tax-base or tax-credit leakage, and the remaining terms represent compliance, timing, audit, and clawback costs.

The chosen route is

$$p^*(c) = \arg \max_{p \in \mathcal{P}_c} \text{NSV}(c, p)$$

subject to stacking and eligibility constraints.

9. Governance, Safety, and Legal Boundary

GoalOS Global Launch is a decision-support, orchestration, evidence, and transaction-control system. It should not impersonate regulated authority. NIST's AI RMF emphasizes governance, context mapping, measurement, and management of AI risk (National Institute of Standards and Technology, 2023). The EU AI Act creates role- and risk-specific obligations for providers, deployers, importers, distributors, and high-risk systems (European Parliament and Council of the European Union, 2024). GoalOS therefore needs explicit role classification and deployment-specific controls.

9.1. No secrets before contract

The public interface should reject confidential, privileged, personal, regulated, classified, credential, vulnerability, and unauthorized third-party information. Sensitive work requires a separately reviewed protected environment, customer agreement, data map, processor assessment, retention policy, security controls, and incident response.

9.2. Professional authority

GoalOS may identify issues, retrieve sources, create comparisons, prepare dossiers, coordinate professionals, test completeness, and record decisions. Licensed and responsible providers supply legal opinions, tax opinions, regulated investment activity, immigration representation, audits, official certifications, and other reserved judgments.

9.3. Fail-closed execution

OpenAI's Agents SDK exposes guardrails that can reject or trip execution before and after tool calls (OpenAI, 2026). GoalOS generalizes this pattern beyond model safety: every material action should have input, tool, policy, evidence, and authority guardrails. An exception should not silently become a manual founder rescue; it should be repriced, conditioned, deferred, escalated, or declined.

10. Recurring Operation: The Sovereign Expansion Office

A launch is episodic; the institution is not. The recurring product monitors:

- programme openings, deadlines, awards, and clawbacks;
- tax-law, treaty, sanctions, export, and regulatory changes;
- entity substance, beneficial ownership, and governance;
- transfer pricing and permanent establishment;
- hiring, payroll, immigration, and worker status;
- investor, cap-table, financing, and covenant events;
- procurement, infrastructure, insurance, and vendor concentration;
- source expiry, transaction authority, and unresolved Proof Debt.

The Office converts the low-frequency formation event into a high-value continuing control plane. Its renewal case is evidence-based: what changed, what was verified, which risks were retired, which opportunities opened, and what must be done next.

11. Recursive Self-Improvement

11.1. Chronicle admission

Suppose mission m_1 yields candidate capability k . Capability k may influence mission m_2 only if:

$$\text{Admit}(k) = 1$$

where admission requires evidence, rights, scope, version, currentness, known failure conditions, reviewer acceptance, and revocation rules.

11.2. Fresh successor test

Let g_1 be the predecessor generation and g_2 the successor using admitted capability. On fresh mission m_2 , define improvement:

$$\Delta(g_2, g_1 \mid m_2) = \alpha\Delta Q + \beta\Delta V + \gamma\Delta T + \delta\Delta C - \lambda\Delta R - \mu\Delta H, \quad (8)$$

where Q is quality, V verified value, T time efficiency, C cost efficiency, R risk, and H human or reviewer burden. Promotion requires a pre-registered positive threshold and no material governance regression.

Succession Law

Reuse is not improvement. A successor earns authority only through a fresh comparison under controlled or explicitly normalized conditions.

12. Illustrative Mission

Consider the objective:

Launch an AI infrastructure company that performs core R&D in Montreal, sells to US and EU enterprises, hires specialists in three countries, preserves founder control, seeks non-dilutive support, and prepares a pilot compute facility.

GoalOS may generate six scenarios:

Table 5: Illustrative scenario frontier.

Scenario	Design logic	Principal Proof Debt
Sovereign Core	Canadian/Quebec root; foreign sales without premature subsidiaries; R&D, capability, and governance concentrated	US procurement, EU data roles, transfer-pricing triggers, facility economics
Capital-Efficient Expansion	Core retained; US subsidiary only when customer or financing threshold passes; EOR for first foreign hires	CCPC control, investor rights, US nexus, EOR-to-entity trigger
Market-Access First	US commercial entity activated early; EU partner route; R&D remains in Quebec	Tax and IP agreements, management substance, incentives interaction
R&D and Capability First	Research and capability centered in Quebec; Horizon and consortium paths explored	Funding eligibility, partner IP, commercialization nexus
Infrastructure First	Project SPV selected through power, land, permit, and offtake auction	Permits, customer commitments, project finance, construction risk
Global Functional Optimum	Every function allocated to its top-scoring jurisdiction	Entity proliferation, coordination, PE, transfer pricing, control, unwind cost

The system should not automatically choose the highest raw score. It should identify the Pareto frontier, expose the cost of complexity, and recommend the least irreversible scenario capable of proving the next consequential uncertainty.

13. Commercialization Architecture

13.1. Product ladder

GoalOS's Autonomous Revenue Institution provides the commercial state machine: create permissioned demand, diagnose the objective, qualify value and fit, price within margin policy, contract, collect payment, reserve

Table 6: Illustrative product ladder; not a current offer or forecast.

Product	Illustrative price	Target gross margin
Global Launch Readiness Scan	Free–US\$10,000	75–90%
Global Architecture Mission	US\$25,000–100,000	70–85%
Full Launch Mission	US\$100,000–500,000	65–80%
Capital & Public-Support Mission	US\$100,000–500,000	70–85%
Infrastructure Launch Mission	US\$250,000–2,000,000+	50–75%
Sovereign Expansion Office	US\$250,000–1,000,000+ annually	75–85%
Protected Launch Intelligence API	Usage-based	85–95%
Strategic or sovereign programme	US\$5,000,000–20,000,000+	40–70%

capacity, execute, prove, renew, expand, and reinvest verified contribution (MONTREAL.AI Research, 2026a).

13.2. Customer value gate

A default qualification law is:

$$V_{12m}^{conservative} \geq 4F,$$

where F is the customer fee. This does not guarantee ROI; it prevents scarce autonomous capacity from being assigned to objectives whose plausible value cannot justify the mission.

13.3. Decacorn and megacorn scenarios

The scale thesis comes from recurring institutional control, not formation fees.

Table 7: Illustrative operating configurations; not valuation forecasts.

Revenue engine	Decacorn operating case	Megacorn candidate case
Sovereign Expansion Offices	$500 \times \text{US\$1M} = \text{US\$500M}$	$3,000 \times \text{US\$1M} = \text{US\$3B}$
Launch and transaction missions	US\$250M	US\$1B
Protected API / network usage	US\$150M	US\$1.5B
Marketplace and validation	US\$100M	US\$0.5B
Infrastructure and strategic programmes	Excluded / optional	US\$1B
Illustrative annual scale	US\$1B	US\$7B

These scenarios indicate what the institution would need to become: a trusted global launch and expansion standard with high recurring revenue, strong retention, a rights-controlled capability network, professional and provider ecosystems, and demonstrated capital-allocation competence.

14. Empirical Proof Programme

The first proof should remain narrow. A credible Commercial Proof Run requires:

A public-safe Receipt Pack should include the paid scope, preregistered acceptance criteria, evidence summary, independent verdict, complete cost denominator, implementation result, limitations, repeat purchase, Chronicle decision, and exact claim being made.

Table 8: Minimum evidence gates for the first external proof programme.

Gate	Minimum
Unrelated paying customers	3
Substantially similar Launch Missions	3
Independently verified customer outcome	1
Public-support, fiscal, financing, or market route actually realized	1
Completed entity, market, talent, capital, or infrastructure activation	1
Separate repeat purchase or expansion	1
Active Sovereign Expansion Office	1
Rights-cleared Chronicle capability	1
Fresh successor improvement	1
Material trust breaches	0

15. Risks and Failure Modes

16. Research Agenda

The architecture suggests a programme of empirical work:

1. **Scenario quality:** compare GoalOS scenario frontiers with independent multidisciplinary teams.
2. **Value of Proof:** estimate which evidence most changes institutional decisions.
3. **Graph efficiency:** measure when parallel AGI Jobs reduce elapsed time and when coordination cost dominates.
4. **Verifier independence:** test context separation, blind evaluation, and evaluator diversity.
5. **Funding-stack realization:** compare gross discovered support with awarded and collected net cash.
6. **Transaction safety:** measure prevented incomplete or contradictory transactions.
7. **Chronicle precision:** estimate false-admission and false-rejection rates for reusable capability.
8. **Successor transfer:** determine whether admitted capability improves fresh missions without governance regression.
9. **Recurring value:** measure whether the Sovereign Expansion Office reduces risk, cost, delay, or missed opportunity.
10. **Institutional externalities:** measure concentration, regulatory, labour, environmental, and infrastructure effects.

17. Conclusion

The world already contains mature tools for incorporation, hiring, equity, capital administration, legal research, public support, procurement, and expansion. GoalOS Global Launch Ω does not need to replace all of them at once. It can coordinate them around the outcome the sponsor actually wants.

Tell GoalOS what you want to build.

Table 9: Principal risks and institutional controls.

Risk	Failure mode	Control
False legal or tax certainty	Strategic output is treated as binding authority	Current official sources, exact facts, professional receipt, expiry, transaction gate
Grant theatre	Unawarded support enters runway or drives the project	Awarded/earned/collected states separated; support excluded from base cash
Shell-company optimization	Tax score dominates substance and operating reality	Business-purpose memo, people/risk map, entity-cost penalty, local substance review
Scenario monoculture	One model anchors on one attractive structure	Six diverse scenarios, independent red team, counterfactual and stop conditions
Metric gaming	System optimizes headline tax, grant size, or number of entities	Verified launch value includes full coordination, risk, delay, and unwind denominator
Professional conflict	Adviser compensation biases the ranking	Compensation disclosure, separated evaluation and referral functions, independent review
Data leakage	Sensitive customer facts enter a public tool	No-secrets gate, protected deployment, minimization, access logs, retention and deletion
Capability overreach	One successful case becomes universal memory	Narrow scope, versioning, known failures, expiry, revocation, fresh transfer
Infrastructure overinvestment	Large assets are funded before demand and permits	Customer commitments, project SPVs, milestone finance, stop and unwind gates
Founder dependence	Every exception returns to one individual	Fixed products, clause matrices, delegated authorities, explicit decline and escalation paths

It turns the world into your launch stack—and the objective into a verified operating institution.

The paper's strongest claim is institutional rather than technological. Better models make research, drafting, analysis, and orchestration dramatically cheaper. The durable opportunity is to convert that capability into a governed system of objectives, permissions, bounded work, evidence, authority, selective memory, recurring control, and productive reinvestment.

Final Institutional Law

No consequential objective, no mission. No exact facts, no jurisdictional or fiscal result. No real dependency, no graph edge. No Evidence Docket, no value claim. No responsible authority, no transaction. No Chronicle admission, no future influence. No fresh transfer, no improvement claim. No verified contribution, no productive reinvestment.

A. Minimum Launch Mission Contract

A production Mission Contract should contain at least:

1. accountable sponsor and decision authority;
2. objective, scope, exclusions, and deadline;
3. functions to optimize;
4. data classes and permitted sources;
5. models, tools, providers, and environments;
6. budget, cost ceiling, and value hypothesis;
7. acceptance tests and evaluator;
8. professional-authority dependencies;
9. Evidence Docket requirements;
10. rights, confidentiality, publication, and reuse;
11. stop, pause, escalation, and rollback rules;
12. Chronicle eligibility and successor-test conditions.

B. Evidence Docket Schema

Docket object	Required content
Objective record	Sponsor, decision, value at stake, deadline, acceptance criteria
Fact register	Verified facts, source, date, owner, confidence, expiry
Assumption ledger	Assumption, sensitivity, evidence needed, decision impact
Claims matrix	Material claim, supporting evidence, contradiction, status
AGI Job ledger	Worker, tools, data, budget, output, failure, verifier
Cost ledger	Models, data, professionals, support, travel, filings, capital, delay
Authority register	Approval, scope, owner, effective date, expiry, revocation
Transaction receipts	Signed agreements, filings, payments, awards, permits, acceptance
Outcome record	What happened, customer value, adverse outcomes, delayed effects
Rights register	Customer rights, platform rights, third-party rights, reuse conditions
Chronicle decision	Admit, admit with limits, repair, reject, revoke, or expire
Successor result	Baseline, fresh mission, metrics, improvement, regressions, verdict

C. Canonical Operating Doctrine

Discover globally.
 Decompose function by function.
 Generate competing architectures.
 Challenge adversarially.
 Resolve locally.
 Govern institutionally.
 Execute only within authority.
 Record the full truth.
 Admit only what passed.
 Revoke when stale.
 Reinvest only verified value.

References

- Anthropic (2026). *Claude Code Documentation: Agent Teams*. <https://code.claude.com/docs/en/agent-teams>. Accessed 2026-07-24.
- Campbell, Donald T. (1979). “Assessing the Impact of Planned Social Change”. In: *Evaluation and Program Planning* 2.1, pp. 67–90.
- Carta (2026). *The End-to-End Suite Connecting Private Capital*. <https://carta.com/>. Official product page; accessed 2026-07-24.
- Coase, Ronald H. (1937). “The Nature of the Firm”. In: *Economica* 4.16, pp. 386–405.
- Deel (Sept. 2025). *Deel-Owned Infrastructure and Entities Available in Over 100 Countries*. <https://www.deel.com/blog/deels-owned-infrastructure-and-entities-are-now-available-in-over-100-countries/>. Accessed 2026-07-24.
- European Parliament and Council of the European Union (2024). *Regulation (EU) 2024/1689 Laying Down Harmonised Rules on Artificial Intelligence*. <https://data.europa.eu/eli/reg/2024/1689/oj>. Official Journal text; accessed 2026-07-24.
- Goodhart, Charles A. E. (1975). “Problems of Monetary Management: The U.K. Experience”. In: *Papers in Monetary Economics*. Reserve Bank of Australia.
- MONTREAL.AI Research (July 2026a). *GoalOS Autonomous Revenue Institution*. Tech. rep. Internal working paper. MONTREAL.AI.
- (July 2026b). *GoalOS Sovereign Productive Intelligence*. Tech. rep. Internal institutional thesis. MONTREAL.AI.
- (July 2026c). *The Capability–Price Lag and GoalOS Economic Architecture*. Tech. rep. Internal working paper. MONTREAL.AI.
- National Institute of Standards and Technology (Jan. 2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. Tech. rep. NIST AI 100-1. Accessed 2026-07-24. U.S. Department of Commerce.
- OpenAI (2026). *OpenAI Agents SDK: Guardrails*. <https://openai.github.io/openai-agents-python/guardrails/>. Accessed 2026-07-24.
- Remote (2026). *Employer of Record Pricing and Global Hiring Coverage*. <https://remote.com/pricing>. Official product page; accessed 2026-07-24.
- Stripe (2026). *Stripe Atlas: Incorporate Your Startup in Delaware*. <https://stripe.com/atlas>. Official product page; accessed 2026-07-24.
- Williamson, Oliver E. (1985). *The Economic Institutions of Capitalism*. New York: Free Press.