

GoalOS Global Jurisdiction, Funding & Fiscal Intelligence Ω

The Verified Operating System for Global Institutional Allocation

A proof-governed institution for discovering, evaluating, financing, and continuously governing the strongest global allocation of markets, talent, R&D, capability, capital, public support, procurement, compute, energy, infrastructure, and value realization.


The world is your institutional resource graph.

Every jurisdiction. Every instrument. Every cost object. Every authority. One evidence-bearing allocation - function by function, mission by mission, and only after the proof passes.

INSTITUTIONAL VISION

From Global Possibility to Admissible Allocation

A single proof-governed chain from mission to productive capacity.



70
nodes

463
routes

43
named

420
authority routes

476
sources

20
transaction
classes

15
control planes

Discover globally. Prove locally. Allocate only after evidence passes.

Counts describe a dated planning and research substrate. They do not imply universal completeness, open funding, legal authority, or realized value.

This paper presents version 2.0 of *GoalOS Global Jurisdiction, Funding & Fiscal Intelligence Ω* (GJFFI Ω), a proof-governed operating institution for discovering, evaluating, financing, and continuously governing the global allocation of institutional functions. The system does not search for one universally “best country.” It decomposes a lawful mission into eight functions - sell, hire, perform R&D, own capability, raise capital, obtain public support, deploy infrastructure, and realize value - and then constructs concentrated, hybrid, and functionally distributed architectures over a versioned global evidence graph.

Version 2.0 extends the architecture with an Apex Allocation layer that treats evidence acquisition itself as an optimization problem; distributionally robust and minimax-regret selection; explicit real-option and reversibility logic; causal allocation graphs; counterexample search; productive-capacity accounting; capital-to-capacity reinvestment rules; a six-case institutional allocation casebook; an event-sourced production architecture; live source-change propagation; policy-as-code with professional authority; trajectory security; and expanded schemas for missions, nodes, instruments, Proof Debt, scenarios, authority, sources, capital allocation, and benchmarking.

The current public reference substrate contains 70 jurisdiction and functional nodes, 463 governed funding, fiscal, procurement, compute, and capital routes, 43 named instruments, 420 official-authority routes, 476 official and mapped source records, 20 controlled transaction categories, 15 legal and institutional control planes, 12 decision-council roles, and eight competing scenario architectures. These counts describe a dated planning substrate, not universal completeness, current programme availability, or realized funding.

The central claim is institutional: abundant candidate intelligence becomes economically consequential only when converted into current evidence, exact eligibility, complete stacking and fiscal interaction, responsible authority, auditable transactions, realized outcomes, selective memory, and productive reinvestment. The paper specifies the architecture and empirical programme required to earn that claim. It does not assert that binding authority, universal coverage, commercial proof, or civilization-scale productive capacity has already been achieved.

Keywords: global institutional allocation; jurisdiction intelligence; public funding; fiscal optimization; public procurement; capital architecture; transfer pricing; infrastructure; agentic AI; proof governance; eligibility; anti-double-dip; robust optimization; real options; productive capacity; institutional memory.

Claim Boundary

This paper presents a testable institutional architecture, commercialization thesis, and research programme. It does not provide binding legal, tax, accounting, investment, employment, regulatory, procurement, funding, or infrastructure advice; universal jurisdictional coverage; guaranteed support; autonomous regulated practice; programme approval; transaction execution; or regulatory immunity. Any binding action requires current official sources, exact facts, applicable qualified professional review, responsible organizational authority, and complete evidence and transaction receipts.

Document Identity and Reading Rule

Author: MONTREAL.AI. **Edition:** Sovereign Fiscal Intelligence Institution, version 2.0.0. **Status:** research monograph, institutional specification, and falsifiable proof programme. **Intelligence snapshot:** 23 July 2026, with selected official-source context reviewed on 24 July 2026. Registry counts, planning scores, value models, price ranges, and scale cases are design objects unless explicitly identified as realized evidence. The governing rule is: *candidate intelligence may search expansively; allocation and transaction authority remain fail-closed.*

The World as an Evidence-Bearing Resource Graph

The scarce advantage is not access to information. It is the ability to coordinate jurisdictions, instruments, cost objects, capital, professional authority, transactions, and productive capacity without allowing plausible output to become institutional truth.


Category. GJFFI Ω is a continuously governed institution that discovers, evaluates, proves, and monitors the best place to sell, hire, perform R&D, own capability, raise capital, obtain public support, deploy infrastructure, and realize value - function by function, mission by mission, and only after the evidence passes.

Apex allocation. The institution optimizes both the architecture and the order in which uncertainty is reduced. It preserves the expected-value, tail-resilient, and minimax-regret champions; searches for counterexamples; and commits only when the marginal value of further proof falls below its cost or a reversible action can safely produce evidence.

Productive capacity. The long-run objective is the governed conversion of verified intelligence into bounded work, accepted proof, authorized action, protected capability, recurring value, capital, and durable productive capacity - including models, people, customers, compute, energy, laboratories, infrastructure, and institutions.

THE EIGHT ALLOCATION FUNCTIONS

One Mission. Eight Functions. Multiple Jurisdictions. One Governed Architecture.



Allocation law. The best jurisdiction for one function need not be the best for another. A customer market cannot erase a fatal funding restriction; a tax advantage cannot replace operating substance; a large grant cannot justify weak rights or negative cash timing; and infrastructure cannot be financed before demand, permits, power, and authority exist.

Contents

1.	Introduction	12
2.	The Fragmented Global-Allocation Stack	13
2.1.	Why a generic AI answer is insufficient	13
3.	Category Definition	14
4.	The Eight Independently Optimized Functions	14
5.	Why the Timing Is Exceptional	15
5.1.	Public support as competitive architecture	15
5.2.	Concentrated global investment	15
5.3.	Minimum tax and incentive interaction	15
5.4.	AI and agentic regulation	15
5.5.	Compute and infrastructure support	15
6.	Mission Constitution	17
7.	Jurisdiction Nodes and the Fifteen-Factor Constitution	17
8.	Institutional Scenario	18
9.	Verified Global Allocation Value	18
10.	Instrument Eligibility and Expected Realizable Support	19
11.	Cost-Object Stacking	19
12.	Proof Debt	19
13.	Transaction Gate	19
14.	Partial Observability, Uncertainty, and Reversibility	20
15.	Value of Proof	20
16.	The Allocation Meta-Problem	22
17.	Whole-Architecture Search	22
18.	Scenario Diversity as a Control	22
19.	Distributionally Robust Selection	23
20.	Real Options and Staged Commitment	24
21.	Value of Information and Proof Allocation	24
22.	Causal Allocation Graph	25
23.	Counterfactual and Counterexample Search	26
24.	Stopping Rules	26
25.	Current Reference Substrate	28
26.	The Seventy-Node Jurisdiction Registry	28

27.	Priority-Node Constitutions	29
28.	The Governed Instrument Registry	30
29.	Evidence-State Machine	31
30.	The Official-Source Registry	32
31.	Source Freshness and Evidence Decay	33
32.	Jurisdiction Activation Gates	33
33.	Eligibility Engine	35
34.	Stacking and Anti-Double-Dip Engine	35
35.	Fiscal Intelligence Engine	36
36.	Capital Stack Engine	37
37.	Entity, Capability, and Transfer-Pricing Engine	37
38.	Procurement and Infrastructure Engine	38
39.	Transaction Engine	38
40.	Monitoring and Update Engine	38
41.	Twelve-Role Decision Council	40
42.	The Ω Deliberation Loop	40
43.	Positive-Upside Search	41
44.	Graph Engineering	41
45.	Institutional Digital Twin	41
46.	Evidence Docket	42
47.	Chronicle and Selective Memory	42
48.	Recursive Improvement	42
49.	Trajectory Security	43
50.	Priority Jurisdiction Archetypes	45
51.	Illustrative Mission: AI Infrastructure Company	45
52.	Illustrative Scenario Frontier	46
53.	Capital and Public-Support Compilation	47
54.	Illustrative Cost-Object Allocation	48
55.	Illustrative Transaction Packet	48
56.	Infrastructure and Productive Capacity	49

57.	Casebook Method	51
58.	Case A: Frontier AI Company	51
58.1.	Coupled decisions	51
58.2.	First proof actions	51
59.	Case B: Advanced-Technology Research Consortium	52
60.	Case C: AI Compute and Energy Campus	52
60.1.	Project architecture	52
61.	Case D: Global Scientific Foundation	53
62.	Case E: Cross-Border Investment Fund	53
63.	Case F: Regional Economic-Development Institution	54
64.	Cross-Case Principles	54
65.	From Fiscal Optimization to Productive Capacity	56
66.	The Compounding Chain	56
67.	Productive-Capacity Classes	57
68.	Capacity Frontier	58
69.	Public Support as Capacity Catalyst	58
70.	Compute, Energy, and Sovereign Infrastructure	58
71.	Capital Allocation Council	59
72.	Network and Civilization-Scale Externalities	59
73.	Compounding Without Overreach	59
74.	Authority Boundary	61
75.	Fifteen Legal and Institutional Control Planes	61
76.	AI and Agentic Governance	62
77.	Privacy and Data Architecture	63
78.	Sanctions, Export Controls, and Investment Screening	63
79.	Professional Authority	63
80.	Conflicts, Marketplace, and Ranking Independence	63
81.	Claims and Publication Controls	64
82.	Deployment Modes	64
83.	Production System Topology	66
84.	Event-Sourced Institutional State	66

85.	The Evidence Graph	66
86.	Source Ingestion and Change Detection	67
87.	Policy-as-Code and Human Authority	67
88.	Agentic Work Fabric	68
89.	Trajectory Security	68
90.	Professional Network and Responsibility Map	69
91.	Deployment Modes	69
92.	Operational Cadence	69
93.	A 1,000-Day Production Path	70
94.	Institutional KPIs	70
95.	Product Architecture	72
96.	Illustrative Product Ladder	72
97.	Customer Value Gate	73
98.	The Recurring Global Expansion Office	73
99.	Business Potential	73
100.	Illustrative Scale Configurations	74
101.	The Institutional Moat	74
102.	The First Commercial Wedge	75
103.	A 1,000-Day Institutional Path	75
104.	GJFFIBench Ω	77
105.	Benchmark Metrics	77
106.	Experimental Protocol	78
107.	Commercial Proof Run	78
108.	Research Agenda	79
109.	Claim Ladder	79
110.	Limitations	81
110.1.	Coverage is not completeness	81
110.2.	Official routes are not open programmes	81
110.3.	Scores are not truth	81
110.4.	Probabilities are difficult to calibrate	81
110.5.	Professional judgment is not fungible	81
110.6.	A browser-local release is not a live intelligence service	81
110.7.	Recursive systems can amplify error	81

110.8.	The benchmark may be gamed	81
110.9.	Commercial scale is not a present capability claim	81
111.	Principal Risks and Controls	82
112.	Conclusion	83
A.	Fifteen-Factor Jurisdiction Constitution	85
B.	Governed Instrument Record	85
C.	Eligibility Rule Set	86
D.	Cost-Object Stacking Ledger	87
E.	Transaction Receipt Set	87
F.	Evidence Docket Schema	87
G.	Source Freshness and Verification States	88
H.	GJFFIBench Mission Record	88
I.	Canonical Operating Doctrine	88
J.	Mission Constitution v2 Schema	89
K.	Jurisdiction Node v2 Schema	89
L.	Instrument v2 Schema	90
M.	Proof Debt Taxonomy	90
N.	Scenario Dossier v2	91
O.	Authority Object	91
P.	Source Snapshot Object	91
Q.	Capital Allocation Receipt	92
R.	GJFFIBench v2 Trial Record	92
S.	Glossary	92
T.	Sovereign Fiscal Operating Doctrine v2	93
	References	93

The Global Allocation Problem

The world contains more capital, public support, talent, infrastructure, law, and strategic opportunity than any institution can evaluate manually. The scarce capability is deciding what fits, what stacks, what survives scrutiny, and what should happen next.

1. Introduction

Global institutional design is usually performed through disconnected conversations. A founder asks counsel where to incorporate; a tax adviser models one structure; a grant consultant searches one programme family; a recruiter proposes a hiring route; an economic-development agency promotes its own jurisdiction; a cloud provider offers credits; a lender evaluates financeability; and a board attempts to reconcile the resulting assumptions. Each actor can be competent while the combined architecture remains incoherent.

The scale of the resource field is extraordinary. Global foreign direct investment reached approximately US\$1.6 trillion in 2025, while the top 20 host economies attracted more than 80% of global FDI and strategic sectors represented 44% of global greenfield project values (United Nations Conference on Trade and Development, 2026b; United Nations Conference on Trade and Development, 2026a). Across only 15 major industrial sectors, the OECD identified US\$108 billion of grants, tax concessions, and below-market finance in 2024 (Organisation for Economic Co-operation and Development, 2026c). The global minimum tax is now a coordinated implementation system spanning a large number of jurisdictions, with a 15% floor for in-scope groups and operational guidance on administration, filing, coordination, and implementation (Organisation for Economic Co-operation and Development, 2026a; Organisation for Economic Co-operation and Development, 2026d). These variables change where investment flows, which activities are financeable, how public support is valued, and whether a nominally attractive structure survives implementation.

The opportunity is not merely to retrieve more information. It is to convert a dynamic global resource field into a mission-specific, evidence-bearing, authority-aware architecture.

**The category is not country comparison.
It is global institutional allocation.**

GJFFI Ω begins with a consequential objective and asks:

Where should every institutional function occur, which instruments should support it, what interactions change the net result, which evidence remains unresolved, and which actions may lawfully proceed?

This paper makes seven contributions:

1. It defines a global allocation mission over eight independently assigned institutional functions.
2. It formalizes a 15-factor jurisdiction constitution and multi-node scenario architecture.
3. It specifies an instrument registry whose records include eligibility, ownership, expenditure, location, IP, stacking, tax, timing, evidence, and clawback conditions.
4. It introduces eligibility, stacking, fiscal-interaction, capital, and transaction engines that preserve the complete economic denominator.
5. It presents a twelve-role agentic decision council, robust scenario frontier, institutional digital twin, and living evidence graph.
6. It defines source currentness, professional authority, legal control, security, and fail-closed transaction rules.
7. It proposes a commercialization architecture and empirical benchmark for determining whether GJFFI Ω creates superior realized value.

2. The Fragmented Global-Allocation Stack

Table 1: The fragmented market and the GJFFI reframing.

Application class	Primary object	Question left unresolved
Tax database	Rates, rules, treaties, summaries	Which architecture fits the exact functions, facts, thresholds, and capital plan?
Grant database	Programme discovery	Is the entity eligible, can the support stack, what is the net cash value, and when will it arrive?
Site-selection tool	Location cost and quality	How do customers, talent, permits, energy, finance, tax, incentives, and reversibility interact?
Formation platform	Entity filing	Should the entity exist, which functions belong there, and what evidence justifies it?
EOR or payroll platform	Hiring without a local entity	When does the complete economic and tax logic support EOR, contractor, branch, or subsidiary?
Capital platform	Equity or debt workflow	Which capital should finance which asset or stage, under what control and evidence gate?
Professional advisory	Bespoke judgment	How are incompatible assumptions reconciled, monitored, and reused without overreach?
Development agency	Promotion of one location	How does that jurisdiction compare with non-promotional alternatives on the complete mission?

The weakness is not that these products are unimportant. Their optimization boundary is narrower than the sponsor's objective. The sponsor bears the transaction and coordination cost of reconciling different providers, definitions, currencies, timelines, incentives, conflicts, and standards (Coase, 1937; Williamson, 1985).

2.1. Why a generic AI answer is insufficient

A frontier model can summarize law, compare tax rates, identify grants, draft an entity map, and generate a persuasive recommendation. That does not create an admissible institutional architecture. Candidate output lacks, unless separately established:

- current official programme and legal status;
- exact entity, ownership, location, expenditure, sector, and threshold facts;
- complete stacking and tax-base interactions;
- responsible legal, tax, fiduciary, audit, procurement, engineering, or regulatory authority;
- evidence that applications, awards, financings, filings, contracts, hiring, and infrastructure actually occurred;
- the full denominator of professional cost, delay, audit risk, clawback, coordination, and unwind;
- rights to reuse the resulting method or customer-derived capability;
- a mechanism for expiry, repair, revocation, or successor comparison.

Better models make candidate cognitive production cheaper. Verified truth, fact-specific eligibility, authority, execution, accountability, rights, and realized outcomes remain scarce.

3. Category Definition

GJFFI Ω is a continuously governed institution that discovers, evaluates, proves, and monitors the best place to sell, hire, perform R&D, own capability, raise capital, obtain public support, deploy infrastructure, and realize value - function by function, mission by mission, and only after the evidence passes.

The system supports lawful companies, nonprofits, research programmes, investment funds, consortia, international expansions, infrastructure projects, industrial facilities, public institutions, economic-development initiatives, and special-purpose vehicles. The same architecture applies because all of them must allocate functions, resources, authority, evidence, and continuing obligations.

4. The Eight Independently Optimized Functions

Let the function set be

$$\mathcal{F} = \{f_{\text{sell}}, f_{\text{hire}}, f_{\text{R\&D}}, f_{\text{capability}}, f_{\text{capital}}, f_{\text{support}}, f_{\text{infrastructure}}, f_{\text{value}}\}.$$

Table 2: The eight allocation functions.

Function	Primary question	Typical evidence
Sell	Where should customers be reached, contracted, billed, procured, and supported?	Demand, channels, public procurement, tax registration, payments, language, service and data roles.
Hire	Where and how should talent be engaged?	Talent depth, fully loaded cost, immigration, worker classification, payroll, benefits, IP assignment, nexus.
Perform R&D	Where should experimental development occur?	People, laboratories, compute, collaboration, project evidence, incentives, eligible expenditure, commercialization nexus.
Own capability	Which entity should own or control reusable capability and IP?	Chain of title, assignments, rights, funding conditions, valuation, DEMPE, substance, export and data controls.
Raise capital	Which entity should raise each class of capital?	Customer cash, investor access, debt capacity, control, securities rules, dilution, guarantees, currency, project finance.
Obtain public support	Which instruments can lawfully support which costs?	Eligibility, cost share, location, ownership, stacking, tax-base effects, cash timing, audit, clawback.
Deploy infrastructure	Where should compute, laboratories, facilities, energy, or industrial assets be built?	Demand, power, land, permits, logistics, finance, offtake, resilience, decommissioning.
Realize value	How should returns, distributions, royalties, exits, reinvestment, or public value occur?	Tax, withholding, treaties, governance, liquidity, impact, productive reinvestment, unwind.

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

5. Why the Timing Is Exceptional

5.1. Public support as competitive architecture

Public support increasingly shapes industrial, technological, and infrastructure competition. The OECD's 2026 firm-level work tracks grants, tax concessions, and below-market finance and reports US\$108 billion across 15 sectors in 2024 (Organisation for Economic Co-operation and Development, 2026c). The implication is not that every institution should chase subsidies. It is that strategic location and financing increasingly depend on public-policy interaction.

5.2. Concentrated global investment

UNCTAD reports that global FDI rose to US\$1.6 trillion in 2025, but the rebound was concentrated in a small number of economies and megaprojects, especially strategic and AI-related digital infrastructure (United Nations Conference on Trade and Development, 2026b). GJFFI Ω must distinguish nominal investment availability from accessible, mission-compatible, financeable capital.

5.3. Minimum tax and incentive interaction

The OECD global minimum tax places a 15% floor under the effective taxation of in-scope multinational income and continues to evolve through implementation tools and safe harbours (Organisation for Economic Co-operation and Development, 2026a; Organisation for Economic Co-operation and Development, 2026b). Tax scores must therefore model scope, substance, timing, qualified domestic minimum top-up taxes, incentives, and exact facts - not merely statutory rates.

5.4. AI and agentic regulation

The EU AI Act applies through staged obligations. Prohibited-practice and AI-literacy rules apply from February 2025, general-purpose AI obligations from August 2025, Article 50 transparency obligations from 2 August 2026, and the 2026 political agreement places key high-risk dates in December 2027 and August 2028 (European Commission, 2026a; European Commission, 2026c; European Commission, 2026b; European Commission, 2026d). Singapore's updated agentic-AI governance framework emphasizes bounded powers, meaningful human accountability, lifecycle controls, and risks from multi-agent and third-party systems (Infocomm Media Development Authority of Singapore, 2026b). Actual roles, uses, data, affected persons, and external actions determine the perimeter.

5.5. Compute and infrastructure support

Public support increasingly reaches compute and infrastructure directly. Canada's AI Compute Access Fund was established with up to C\$300 million; the first call is closed, and 44 supported projects representing C\$66 million were announced in May 2026 (Innovation, Science and Economic Development Canada, 2026a; Innovation, Science and Economic Development Canada, 2026d). Canada also maintains separate compute-infrastructure routes, including the AI Compute Challenge and sovereign-supercomputing initiatives (Innovation, Science and Economic Development Canada, 2026b; Innovation, Science and Economic Development Canada, 2026c). The European Innovation Council's 2026 programme exceeds EUR 1.4 billion across schemes with materially different grant, investment, eligibility, and application structures (European Innovation Council, 2026c; European Innovation Council, 2026a; European Innovation Council, 2026b). The correct unit is a dated, fact-specific, evidence-bearing route.

Formal Allocation Model

The institution must distinguish high expected upside from admissible action.
Value, evidence, authority, and reversibility are separate variables.

6. Mission Constitution

A global allocation mission is represented as

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

where O is the consequential objective; S the accountable sponsor and stakeholder set; B approved budget and capital; τ time horizon and deadlines; R risk, control, and reversibility posture; A the authority model; $\mathcal{F}$ the functions to allocate; Π policy and regulatory constraints; and E the initial evidence state.

A production Mission Constitution should freeze at least:

1. sponsor, decision owner, affected stakeholders, and delegated authority;
2. objective, scope, exclusions, deadline, value at stake, and success criteria;
3. entity facts, ownership, control, revenue, headcount, locations, sector, and regulated activities;
4. functions to allocate and the maximum permitted architecture complexity;
5. capital available, matching capital, dilution tolerance, liquidity, and runway;
6. data classes, permitted sources, models, tools, environments, and confidentiality;
7. evidence threshold, evaluator, professional dependencies, and expiry;
8. prohibited jurisdictions, sanctions, export, data, sovereignty, and policy constraints;
9. stop, pause, escalation, rollback, and publication rules;
10. Chronicle eligibility and successor-test conditions.

Constitution Law

No optimization may silently alter the sponsor's objective, authority, prohibited actions, evidence threshold, or risk posture. A material change creates a new mission version and a new decision record.

7. Jurisdiction Nodes and the Fifteen-Factor Constitution

Let C_t be the set of candidate jurisdiction and functional nodes at time t . For node c and function f , define the evidence-adjusted score

$$JScore(c, f \mid M, t) = \sum_{k=1}^{15} w_{f,k}(M) z_{c,k,t} q_{c,k,t} - P_{c,f}(M, t),$$

where $w_{f,k}$ is the mission-specific factor weight, $z_{c,k,t}$ the normalized planning score, $q_{c,k,t} \in [0, 1]$ evidence quality and currentness, and $P_{c,f}$ a penalty for blockers, contradictions, coordination, concentration, or Proof Debt.

The fifteen factors are:

1. customer access;
2. talent;
3. R&D incentives;
4. commercialization incentives;
5. effective tax;
6. capability and IP treatment;
7. financing;
8. procurement;

9. AI and data regulation;
10. export controls and investment screening;
11. energy and infrastructure;
12. cost;
13. reversibility;
14. treaty network;
15. political and legal stability.

Scores are not timeless facts. They are dated priors for directing evidence acquisition. A high score cannot erase a fatal ownership restriction, prohibited end use, unavailable power connection, missing permit, or incompatible grant covenant.

8. Institutional Scenario

A scenario s is represented as

$$s = (r, J, L, K, P, T, G, Q),$$

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

The system should generate a frontier rather than a single answer. At minimum:

- concentrated sovereign core;
- capital-efficient expansion;
- market-access first;
- R&D and capability first;
- infrastructure first;
- minimum regulatory surface;
- minimum dilution;
- global functional optimum.

A candidate is feasible only when hard constraints hold. Among feasible candidates, Pareto dominance should preserve architectures that trade speed, value, evidence, control, complexity, and resilience differently.

9. Verified Global Allocation Value

The objective is not minimum tax, maximum gross grant, or maximum number of high-scoring nodes. Define

$$\begin{aligned} VGAV(s) = & V_{market}(s) + V_{capital}(s) + V_{support}(s) + V_{talent}(s) + V_{infra}(s) \\ & + V_{tax}(s) + V_{capability}(s) + V_{option}(s) \\ & - C_{formation}(s) - C_{operating}(s) - C_{coordination}(s) - C_{compliance}(s) \\ & - C_{dilution}(s) - C_{delay}(s) - C_{risk}(s) - C_{concentration}(s) - C_{unwind}(s). \end{aligned}$$

The recommended scenario is

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

$Gate(s) = 1$ only when every material evidence and authority requirement required at the current decision stage has passed. A scenario may have exceptional expected value and remain inadmissible.

10. Instrument Eligibility and Expected Realizable Support

For instrument p , entity and project facts x , and time t , define eligibility

$$Elig(p, x, t) \in \{0, 1, ?\},$$

where ? means unresolved. The rule engine evaluates residence, ownership, control, size, revenue, headcount, sector, project location, technology readiness, eligible spend, matching capital, consortium requirements, prior assistance, IP conditions, deadline, and runway.

Headline award size is not value. Let

$$ERS(p, x) = Pr(award \mid x, E) \cdot Cash_{net}(p, x) \\ - C_{application} - C_{compliance} - C_{delay} - R_{audit} - R_{clawback} - R_{strategic\ restriction}.$$

A closed programme can still be valuable as a monitored route, but its current expected cash value is zero until a new call exists and eligibility is re-established.

11. Cost-Object Stacking

For cost object c and candidate instruments $p \in \mathcal{P}_c$, define net support value

$$NSV(c, p) = G(c, p) - A(c, p) - T(c, p) - C_{professional}(c, p) \\ - C_{compliance}(c, p) - C_{delay}(c, p) - R_{audit}(c, p) - R_{clawback}(c, p),$$

where G is gross support, A assistance interaction, and T tax-base, tax-credit, minimum-tax, or accounting leakage.

The selected allocation is a constrained optimization over cost objects, programmes, tax bases, assistance caps, timing, and evidence. Each dollar of expenditure may be allocated only once unless the governing rules expressly permit an interaction.

12. Proof Debt

For scenario s , let

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

be unresolved claims, missing evidence, ambiguous authority, stale sources, regulatory dependencies, funding uncertainty, or transaction conditions. The weighted Proof Debt is

$$PD(s) = \sum_i m_i u_i e_i,$$

where m_i is materiality, u_i uncertainty, and e_i exposure if the claim is wrong.

Proof Debt should be surfaced, priced, assigned, and retired. It must not be compressed into a persuasive narrative or hidden inside one confidence score.

13. Transaction Gate

For material transaction t , define the required receipt set

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

Then

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

Transaction Law

No entity, IP transfer, licence, financing, option grant, intercompany agreement, regulated filing, major procurement, or infrastructure SPV proceeds merely because a model recommends it. If $TG(t) = 0$, execution remains blocked.

14. Partial Observability, Uncertainty, and Reversibility

Global allocation is a partially observable, path-dependent control problem. Legal rules can be clear while administrator practice remains uncertain; programmes can exist while budgets are exhausted; a tax position can depend on future people and decisions; infrastructure can fail through delay rather than technical impossibility.

GJFFI Ω should therefore combine:

- expected-value analysis;
- robust optimization over parameter ranges (Ben-Tal et al., 2009);
- downside and tail measures such as P10 and CVaR (Rockafellar and Uryasev, 2000);
- minimax regret;
- real-option value and staged commitment (Dixit and Pindyck, 1994);
- explicit switch conditions and reversible experiments;
- residual uncertainty that is not forced into a probability.

15. Value of Proof

Proof resources should be concentrated where they can change a consequential decision. For evidence action a , define

$$VoP(a) = \frac{\mathbb{E}[\Delta DecisionValue \mid a]}{Cost(a) + Delay(a) + Risk(a)}.$$

High-value proof may include an administrator confirmation, customer commitment, legal opinion, power study, tax modelling, site permit, supplier quotation, valuation, or controlled pilot. The system should stop researching low-impact details while a small number of decision-changing uncertainties remain unresolved.

Meta-Optimization Loop

GJFFI Ω optimizes not only the architecture, but also the order in which uncertainty is reduced. It asks: which question, source, professional opinion, experiment, or customer commitment most efficiently collapses the scenario frontier without creating irreversible exposure?

Apex Allocation and Opportunity Search

The institution does not merely rank countries. It searches the global possibility space, preserves non-dominated architectures, prices uncertainty, and directs proof toward the decisions where it can create the greatest verified upside.

16. The Allocation Meta-Problem

The ordinary optimization problem asks which jurisdiction, instrument, or capital source scores highest. The institutional problem is more demanding. The system must determine:

1. which functions should be separated or co-located;
2. which architectures are feasible under law, funding, control, timing, and capacity constraints;
3. which evidence would most change the preferred architecture;
4. which commitments should remain reversible;
5. when the search should stop and authorized execution should begin;
6. when a previously accepted architecture should be reopened.

This produces a three-level optimization:

$$\max_{\pi_{proof}} \mathbb{E} \left[\max_{s \in \mathcal{S}(E\pi)} VGAV(s \mid E\pi) \right] - C(\pi_{proof}),$$

subject to hard legal, authority, budget, privacy, sanctions, and transaction constraints. The outer policy π_{proof} determines what to learn; the inner search constructs scenarios; the transaction layer decides what may actually happen.

Apex Allocation Doctrine

Search may be extremely broad. Evidence acquisition must be targeted. Authority must remain narrow. Irreversible commitments should occur only after the scenario frontier has collapsed sufficiently around an admissible architecture.

17. Whole-Architecture Search

Let $x_{f,c} \in \{0, 1\}$ indicate that function f is assigned to jurisdiction node c . A simplified whole-architecture search is:

$$\begin{aligned} \max_x \quad & \sum_{f \in \mathcal{F}} \sum_{c \in \mathcal{C}} x_{f,c} JScore(c, f) + \sum_{f,g} \sum_{c,d} x_{f,c} x_{g,d} Synergy_{fgcd} \\ & - \sum_{f,g} \sum_{c,d} x_{f,c} x_{g,d} Friction_{fgcd} - Complexity(x) - ProofDebt(x), \end{aligned}$$

subject to:

$$\sum_c x_{f,c} = 1 \quad \forall f,$$

plus mission-specific constraints for ownership, residence, data, export, investment screening, capital, infrastructure, substance, staffing, and professional authority.

The synergy term captures combinations such as R&D near specialist talent and compute, public support matched to eligible costs, or customer access aligned with procurement. The friction term captures transfer pricing, permanent establishment, cross-border data, duplicated compliance, entity proliferation, banking, currency, and coordination.

18. Scenario Diversity as a Control


A single optimizer can converge prematurely on an attractive but fragile structure. GJFFI Ω therefore constructs diverse institutional archetypes before ranking:

- **Sovereign core:** concentrate control, capability, and evidence around one substantive root;

- **Capital-efficient expansion:** add foreign nodes only when a customer, funding, hiring, or financing trigger passes;
- **Market-access first:** privilege revenue, procurement, and distribution;
- **R&D and capability first:** privilege talent, research support, chain of title, and capability compounding;
- **Infrastructure first:** privilege power, land, compute, permits, offtake, and project finance;
- **Minimum regulatory surface:** prefer fewer regulated roles, lower data exposure, and reversible commercial routes;
- **Minimum dilution:** prioritize customer cash, earned credits, awarded support, and contract-backed finance;
- **Global functional optimum:** assign each function to its local optimum, then price coordination and unwind.

A scenario frontier preserves architectures that are not dominated simultaneously on verified value, evidence, speed, control, resilience, complexity, and reversibility.

Apex Scenario Frontier: Value, Proof Debt, Complexity, and Resilience



Bubble size represents architecture complexity. Illustrative values; not a customer recommendation.

Figure 1: Illustrative scenario frontier. High projected value is not sufficient when Proof Debt or architecture complexity remains material.

19. Distributionally Robust Selection

Estimated distributions are often wrong precisely where decisions are consequential. Instead of optimizing only under one forecast distribution P_0 , the institution can use an ambiguity set $\mathcal{P}$ around the available evidence:

$$DRValue(s) = \inf_{P \in \mathcal{P}} \mathbb{E}_P[VGAV(s)].$$

The robust champion need not have the highest expected value. It is the architecture whose advantage survives plausible misspecification of grant probability, tax leakage, customer timing, power availability, regulatory cost, financing terms, or implementation delay.

A complementary minimax-regret rule is:

$$s^{MR} = \arg \min_s \max_{\omega \in \Omega} [VGAV(s^*(\omega), \omega) - VGAV(s, \omega)] .$$




Figure 2: Illustrative risk-adjusted outcome distributions. The highest expected-value architecture need not minimize tail loss or worst-case regret.

Recommended Decision Surface

The system should present at least three answers: the expected-value champion, the tail-resilient champion, and the minimax-regret champion. When they differ, the disagreement is itself decision-changing evidence.

20. Real Options and Staged Commitment

An institutional design can be valuable because it preserves options. A remote-sales route, EOR hire, pilot project, non-binding site reservation, conditional grant application, or customer-funded proof mission may reveal critical information without committing to a subsidiary, IP migration, large facility, or financing round.

For staged action a , define:

$$ROV(a) = \mathbb{E}[V_{future\ options} | a] - C_a - Irreversibility_a.$$

The preferred first move is often not the terminal architecture. It is the least irreversible action capable of proving the next consequential uncertainty.

21. Value of Information and Proof Allocation

For uncertain variable θ and evidence action a , the expected value of sample information is:

$$EVSI(a) = \mathbb{E}_{y \sim p(y|a)} \left[\max_s \mathbb{E}[VGAV(s) | y] \right] - \max_s \mathbb{E}[VGAV(s)].$$

The net value of proof is:

$$NVOP(a) = EVSI(a) - Cost(a) - Delay(a) - Exposure(a).$$

High-value proof actions may include:

- a written programme-administrator interpretation;
- a customer letter of intent or prepaid pilot;
- a legal or tax opinion on a narrow issue;
- a transfer-pricing comparability study;
- a power-availability or interconnection study;
- a permit pre-consultation;
- a lender or investor indication;
- a controlled employment or market-entry pilot;
- a source refresh that determines whether an instrument is actually open.




Figure 3: Illustrative proof-allocation curve. The institution should fund evidence that changes a consequential decision, not research that merely increases narrative detail.

22. Causal Allocation Graph

Correlation is insufficient for policy, incentive, tax, or infrastructure decisions. GJFFI Ω should represent causal hypotheses explicitly:

Jurisdiction $\rightarrow$ *Talent* $\rightarrow$ *Execution* $\rightarrow$ *Revenue*,

Instrument $\rightarrow$ *CashTiming* $\rightarrow$ *Runway* $\rightarrow$ *Dilution*,

EntitySubstance $\rightarrow$ *TaxPosition* $\rightarrow$ *AuditRisk*,

PowerAvailability $\rightarrow$ *CommissioningDate* $\rightarrow$ *CustomerValue*.

Each causal edge carries: mechanism, evidence, confounders, date, owner, uncertainty, and conditions under which the edge should be revoked. Scenario simulation may use causal models where justified, but the system must label structural assumptions rather than hide them in a score.

23. Counterfactual and Counterexample Search

Before admission, the system actively searches for architectures that would overturn the preferred recommendation:

- What if the headline grant is unavailable or delayed by twelve months?
- What if the parent loses a required ownership status after financing?
- What if the first foreign employee creates a permanent establishment?
- What if cross-border data transfer is prohibited for the use case?
- What if the preferred site cannot secure power or permits?
- What if a strategic investor creates foreign-investment review exposure?
- What if an incentive requires IP, employment, or production commitments that reduce future option value?

A scenario survives only if its decision logic remains coherent under material counterexamples or if explicit switch conditions exist.

24. Stopping Rules

Unlimited analysis can destroy value through delay. Search should stop when:

1. the preferred scenario is admissible;
2. no unresolved item exceeds the materiality threshold;
3. the expected value of additional proof is below its cost and delay;
4. the decision is reversible enough to learn through action;
5. the responsible authority accepts the residual uncertainty;
6. the transaction receipts required for the current stage are complete.

Allocation Stopping Law

Do not demand certainty where a reversible experiment can produce evidence. Do not accept experimentation where the action is irreversible, regulated, or capable of producing material harm. The burden of proof rises with irreversibility, externality, and authority.

The Governed Intelligence Substrate

Scale is useful only when every record carries provenance, state, expiry, ownership, and a clear boundary between discovery and transaction readiness.

25. Current Reference Substrate

The current public reference dataset is dated 23 July 2026. It contains 70 jurisdiction and functional nodes, 463 governed routes, 43 named instruments, 420 official-authority discovery routes, 476 official and mapped sources, 20 controlled transaction categories, and 15 legal control planes.




Figure 4: Reference-substrate composition. The counts describe governed planning records, not universal completeness or presently available cash.

Registry-Class Disclosure

Named instruments and official-authority routes are different objects. A named instrument carries programme-specific terms. An official-authority route identifies the responsible channel through which current terms must be discovered and verified. Both are useful; neither may be misrepresented.

26. The Seventy-Node Jurisdiction Registry

The registry contains national, subnational, regional, and functional nodes. Subnational nodes matter because taxes, incentives, employment, procurement, infrastructure, and permitting often differ materially within one country. Regional nodes matter because supranational law, funding, procurement, or market access can operate above the member-state level. Functional nodes represent cross-border or institutional routes that are not reducible to one sovereign territory.




Figure 5: Illustrative distribution of the seventy-node reference registry.

Every node should have:

- stable identity, region, tier, intended role, and version;
- the fifteen factor scores and evidence discounts;
- best uses and explicit watch items;
- relevant legal, fiscal, funding, procurement, talent, and infrastructure source objects;
- a named internal owner and required external professional roles;
- next review, evidence expiry, and activation status;
- entry, stop, rollback, and exit conditions.

27. Priority-Node Constitutions

A priority node needs more than a score. It requires a compact operating constitution that answers:

1. What role can this node play in the mission?
2. Which laws, programmes, procurement routes, and authorities govern it?
3. Which entity, ownership, substance, project-location, and hiring facts activate obligations or eligibility?
4. What data, capability, sanctions, export, investment-screening, and sector constraints apply?
5. What is the complete formation, operating, professional, compliance, and unwind cost?
6. Which sources are current, which remain unresolved, and who is accountable for review?
7. What event should activate, pause, repair, or exit the node?




Figure 6: Illustrative priority-node factor matrix. These are planning scores for evidence acquisition, not legal, fiscal, or investment conclusions.

The heatmap demonstrates why a single rank is insufficient. A node can be excellent for customer access and financing while weak for cost or regulatory simplicity. Another can be powerful for R&D and public support while unsuitable for early commercial contracting. The institution composes these differences rather than averaging them away.

28. The Governed Instrument Registry

The instrument registry covers public support, tax and fiscal incentives, procurement revenue, compute and infrastructure support, export and trade finance, guarantees, debt, equity, and project finance.

Governed Route Registry by Instrument Family




Figure 7: Route families in the current reference registry. Official-authority routes require live terms verification before reliance.

A production record should contain:

- instrument, jurisdiction, authority, type, and official source;
- opening, deadline, decision date, cash date, budget, award, and cost share;
- eligible entity, ownership, size, sector, consortium, and establishment conditions;
- eligible expenditure, project location, technology readiness, and start-date conditions;
- IP, data, publication, commercialization, localization, and change-of-control terms;
- procurement path, application process, evaluation, evidence, and responsible owner;
- stacking, assistance intensity, tax-base, accounting, withholding, and minimum-tax interactions;
- reporting, audit, clawback, security, retention, and exit conditions;
- status, verification date, confidence, open gates, next action, and expiry.

29. Evidence-State Machine

Instrument and node records should move through explicit states rather than one binary “verified” field.




Figure 8: Evidence-state progression. A record can regress when the source, facts, professional conclusion, or governing conditions change.

The canonical states are:

- Discovered** A candidate route or issue exists.
- Sourced** The official authority and primary source are identified.
- Screened** Preliminary facts have been tested; unresolved gates are explicit.
- Verified** Current terms and exact facts support the route at the stated scope.
- Reviewed** Applicable qualified professionals have examined the material issues.
- Authorized** The responsible organization has approved the bounded action.
- Realized** The external result, cash, activation, risk reduction, or value occurred.
- Audited** The evidence survived the applicable audit, reporting, or post-completion review.

A route can become stale, repair-required, rejected, expired, or revoked. Evidence state is a governed object, not marketing language.

30. The Official-Source Registry

The source registry should distinguish:

- constitutions, statutes, regulations, official journals, and case law;
- regulator, tax-authority, procurement, programme, and development-agency materials;
- official application guides, budgets, calls, forms, FAQs, and contribution agreements;
- international standards and intergovernmental guidance;
- professionally authored conclusions and their precise scope;
- non-authoritative commentary, market data, and secondary research.

Every source object requires a retrieval date, content date, jurisdiction, category, authority level, language, relevant claims, verification notes, and invalidation triggers. Archived snapshots matter because a future reviewer must be able to reconstruct which terms governed the decision.

31. Source Freshness and Evidence Decay

Let source quality at time t be

$$q_i(t) = q_i(t_0)e^{-\lambda_i(t-t_0)} \prod_j I_{ij},$$

where λ_i reflects the change rate and I_{ij} are invalidation factors such as a new budget, amendment, election, court decision, administrator notice, programme closure, or changed customer facts.

Different objects decay differently. A treaty may remain stable for years; an application deadline can expire in days; an administrator's guidance can change after a budget; a legal opinion can become inapplicable when ownership, people, data, or transactions change.




Figure 9: Illustrative evidence-decay profiles. Deadlines and programme terms can expire rapidly, while statutes and treaties usually decay more slowly but remain sensitive to amendments and fact changes.

Currentness Law

A current URL is not sufficient. A material claim is current only when the official source, applicable version, exact facts, interpretation, authority, and transaction context remain valid together.

32. Jurisdiction Activation Gates

No node should become operational merely because it scores well. Each active node must pass:

1. **Demand gate:** the function has a real economic or institutional reason to occur there.
2. **Law gate:** applicable roles, obligations, restrictions, licences, and reserved judgments are mapped.
3. **Architecture gate:** entity, data, people, IP, systems, and control design can satisfy the requirements.
4. **Economics gate:** complete tax, incentive, professional, coordination, delay, risk, and unwind economics are acceptable.
5. **Local-capability gate:** responsible people, providers, banking, payroll, procurement, or infrastructure can operate.
6. **Evidence gate:** current sources and required professional receipts exist.
7. **Exit gate:** stop conditions, reversibility, termination, data, people, grant, and asset unwind are defined.

Eligibility, Fiscal, Capital, and Transaction Engines

A route becomes economically meaningful only when exact facts pass, every interaction is modelled, cash timing is real, and authority is complete.

33. Eligibility Engine

The eligibility engine compiles programme rules into explicit tests. It should not rely on one opaque model judgment. For each instrument, the engine produces:

- facts satisfied;
- facts failed;
- facts unresolved;
- sources and retrieval dates;
- evidence required;
- professional interpretation required;
- decision deadline and runway interaction;
- next action and stop condition.

At minimum, it tests:

1. entity residence, establishment, registration, and project location;
2. ownership, beneficial ownership, foreign control, and change-of-control conditions;
3. employee count, revenue, balance sheet, age, stage, and group aggregation;
4. sector, technology, prohibited activity, and regulated-industry status;
5. technology-readiness level, novelty, collaboration, and commercialization route;
6. eligible expenditure, dates, payroll, compute, equipment, contractors, travel, and overhead;
7. matching capital, minimum project size, cost share, liquidity, and runway;
8. IP ownership, publication, licence, data, open-source, and localization requirements;
9. consortium, academic, customer, supplier, regional, or public-sector partner requirements;
10. prior assistance, cumulation, tax-base interaction, aid intensity, and double-dip restrictions;
11. deadline, project start, decision timing, and reimbursement timing;
12. sanctions, export controls, procurement debarment, security, and beneficial-owner screening.

Conservative Screen

A browser-local engine can perform preliminary rule matching. It should use three outcomes - eligible, ineligible, unresolved - and route material ambiguity to current official sources and qualified review. “Likely eligible” is not approval.

34. Stacking and Anti-Double-Dip Engine

For each dollar of expenditure, the system should represent:

Cost object → Proposed route → Alternative route → Interaction → Net value → Chosen allocation.

The model must account for:

- federal, provincial, state, municipal, regional, and supranational interactions;
- tax-credit basis reductions and grant taxability;
- assistance-intensity ceilings;

- expenditure or project exclusions;
- refundable versus non-refundable timing;
- matching requirements and restricted cash;
- professional and compliance cost;
- delayed reimbursement and working-capital cost;
- audit, clawback, covenants, and location or headcount commitments;
- minimum-tax and safe-harbour effects for in-scope groups.




Figure 10: Illustrative conversion of headline support into net realizable support value.

No Double-Dip Law

Every cost object has one canonical ledger. No cost is allocated to two routes unless the governing rules expressly permit the interaction and the tax, assistance, audit, and evidence effects are recorded.

35. Fiscal Intelligence Engine

The fiscal engine must model the institution rather than one tax rate. It covers:

- corporate income tax, losses, credits, incentives, timing, and effective cash tax;
- withholding, treaties, beneficial ownership, anti-abuse, and permanent establishment;
- indirect tax, customs, payroll, property, digital, and sector levies;
- transfer pricing, intercompany agreements, DEMPE, services, financing, and cost sharing;
- controlled foreign company rules, foreign affiliates, exit tax, and migration;
- Pillar Two scope, qualified domestic minimum top-up tax, safe harbours, and top-up interaction;
- accounting, recognition, capitalization, deferred tax, and currency;
- audit probability, controversy, documentation, penalties, and professional burden.

The correct output is a range of after-tax, after-compliance, after-delay cash flows tied to exact facts. A nominal rate has low decision value when the entity lacks substance, treaty eligibility, eligible expenditure, or cash to bridge reimbursement timing.

Substance Before Rate

A low headline rate does not compensate for weak business purpose, absent people and decision-making, artificial risk allocation, ineligible treaty access, funding incompatibility, excessive coordination, or high unwind cost.

36. Capital Stack Engine

The capital engine sequences resources by evidence and economic cost rather than prestige. A disciplined default prior is:

Customer Cash → Earned Credits → Awarded Grants → Approved Debt
→ Equity → Strategic Capital → Project Finance.

The sequence is not universal. It controls unnecessary dilution, speculative runway, and mismatch between capital and asset.

Proof-Gated Capital Sequence



Each resource remains separated as available, eligible, applied, approved, committed, collected, restricted, audited, and released.

Figure 11: Illustrative proof-gated capital sequence.

Each resource remains separately classified as available, eligible, applied, approved, committed, collected, restricted, audited, and repaid or released.

Runway Law

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

37. Entity, Capability, and Transfer-Pricing Engine

The engine evaluates whether an entity is operationally justified. It models:

- business purpose and decision rights;
- functions, assets, risks, people, contracts, systems, and capital;
- customer and supplier relationships;
- capability creation, legal ownership, chain of title, funding conditions, and rights;
- intercompany services, licences, cost sharing, financing, guarantees, and stewardship;
- valuation, transfer pricing, withholding, indirect tax, and minimum-tax interaction;
- local substance, directors, management, payroll, banking, insurance, records, and annual cost;

- foreign-investment, export, sanctions, data, and sector restrictions;
- formation trigger, entity-to-EOR trigger, restructuring trigger, and exit route.

A node is not “optimized” when it reduces one tax line while increasing entity count, management burden, grant incompatibility, audit risk, financing friction, and future unwind.

38. Procurement and Infrastructure Engine

The procurement engine can coordinate legal, tax, accounting, payroll, EOR, banking, insurance, cloud, compute, laboratories, real estate, energy, logistics, security, engineering, construction, and validation providers. It separates provider discovery from recommendation so that compensation cannot silently affect scenario ranking.

The infrastructure engine connects:

Demand → Site → Power → Permits → Capital → Procurement → Commissioning → Output.

It should model land, power, water, connectivity, compute, equipment, labour, logistics, security, permitting, community, insurance, vendor concentration, offtake, project finance, milestones, decommissioning, and environmental or social obligations. Large irreversible assets require customer evidence, permits, lender conditions, and stage gates.

39. Transaction Engine

No entity formation, capability transfer, licence, financing, option grant, intercompany agreement, funding application, procurement commitment, regulated filing, or infrastructure SPV reaches execution-ready state without its applicable receipt set.

The transaction engine should:

1. classify the action, jurisdiction, parties, value, data, rights, and external effect;
2. identify required professional and organizational authorities;
3. open the nine-receipt packet;
4. track dependencies, drafts, reviews, signatures, filings, payments, and implementation;
5. block execution on stale, contradictory, missing, or unauthorized receipts;
6. preserve a tamper-evident Chronicle record;
7. schedule post-completion, covenant, audit, and revocation review.

No Shell Theatre

No entity, licence, invoice, director, risk allocation, or intercompany agreement may be used to disguise the institution’s actual functions, ownership, control, people, data, rights, counterparties, or money flows.

40. Monitoring and Update Engine

The monitoring loop is:

Observe → Detect → Classify → Validate → Rerun → Supersede.

It should support weekly source refresh, monthly funding and fiscal committee, quarterly jurisdiction auction, annual substance and transfer-pricing review, and event-triggered alerts for programme openings, deadlines, budgets, laws, guidance, sanctions, ownership, financing, hiring, customers, infrastructure, evidence expiry, and transaction conditions.

A source update does not automatically change the institution. It creates a candidate claim. The affected record, scenario, transaction, and Chronicle object must be identified; the material effect must be tested; qualified review must occur where necessary; and the change must be versioned.

Agentic Decision Architecture

The system may search more broadly than any committee while remaining subordinate to exact facts, bounded tools, independent proof, responsible authority, and revocable memory.

41. Twelve-Role Decision Council

GJFFI Ω decomposes the mission across specialist decision roles. In a production system these roles may be implemented by models, deterministic engines, humans, professionals, or hybrid workflows.

Table 3: Twelve-role GJFFI decision council.

Role	Institutional mandate
Mission Compiler	Convert the objective, facts, authority, value, constraints, evidence threshold, and stop rules into a frozen Mission Constitution.
Jurisdiction Cartographer	Score and allocate functions across national, subnational, regional, and functional nodes.
Funding Hunter	Discover named programmes and official-authority routes; compile timing, eligibility, evidence, and next actions.
Fiscal Synthesizer	Model tax, incentives, assistance interaction, accounting, cash timing, minimum tax, controversy, and net value.
Entity and Capability Architect	Design the minimum defensible entity, ownership, capability, IP, and intercompany architecture.
Regulatory Sentinel	Classify AI, privacy, cyber, sector, employment, procurement, sanctions, export, and investment-screening exposure.
Capital Engineer	Sequence customer cash, credits, grants, debt, equity, strategic capital, guarantees, and project finance.
Infrastructure Strategist	Evaluate power, compute, land, facilities, logistics, permits, vendors, resilience, finance, and commissioning.
Adversarial Auditor	Attack hidden assumptions, stale evidence, concentration, clawbacks, conflicts, irreversible commitments, and reward gaming.
Scenario Synthesizer	Generate non-dominated whole-institution architectures, explain trade-offs, and preserve switch conditions.
Robustness Laboratory	Simulate parameter and regime uncertainty; measure P10, CVaR, regret, resilience, and scenario failure.
Chronicle Governor	Control evidence state, rights, admission, expiry, revocation, lineage, and fresh-successor comparison.

42. The Ω Deliberation Loop

The canonical loop is:

SENSE $\rightarrow$ CONSTITUTE $\rightarrow$ GENERATE $\rightarrow$ SIMULATE
 $\rightarrow$ CHALLENGE $\rightarrow$ PROVE $\rightarrow$ ADMIT $\rightarrow$ LEARN.

1. **Sense:** collect only authorized, decision-changing facts and source changes.
2. **Constitute:** freeze mission, authority, risk, evidence, and stop conditions.
3. **Generate:** construct diverse whole-institution architectures, not cosmetic variants.
4. **Simulate:** estimate value, cash, time, tail loss, complexity, and regret under uncertainty.
5. **Challenge:** search for legal, fiscal, funding, data, security, capital, infrastructure, and execution counterexamples.
6. **Prove:** acquire high-value evidence and independent review.
7. **Admit:** permit only bounded actions or capabilities whose current gate has passed.

8. **Learn:** preserve outcome, failure, rights, limitations, and successor evidence without silently rewriting authority.

43. Positive-Upside Search

An advanced system should search for extraordinary, non-obvious configurations without confusing novelty with authority. Candidate generation may explore:

- unusual function allocations across jurisdictions;
- customer-funded or procurement-led launch paths;
- combinations of refundable credits, grants, guarantees, strategic capital, and project finance;
- phased market entry and delayed entity activation;
- consortium, university, sovereign, supplier, or infrastructure partnerships;
- customer commitments that unlock financing and reduce subsidy dependence;
- structures that preserve ownership while increasing productive capacity;
- option-like experiments that resolve consequential uncertainty before irreversible commitment.

Each candidate must still disclose complete costs, dependencies, conflicts, uncertainty, authority, failure modes, and unwind.

Opportunity Search Law

More computation should expand the candidate frontier and deepen adversarial search. It may not lower the evidence threshold, bypass professional authority, or convert simulated upside into realized value.

44. Graph Engineering

GJFFI Ω distinguishes five edge types:

1. **Data edge:** a downstream node consumes an upstream artifact.
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 capital, time, data, tools, suppliers, or state.
5. **Governance edge:** a result cannot influence later missions without Chronicle admission.

This prevents two common errors: treating every narrative sequence as a computational dependency, and treating every technical completion as institutional permission.

The living graph connects:

Mission → Functions → Jurisdictions → Instruments → Claims
→ Evidence → Authority → Transactions → Outcomes → Chronicle.

45. Institutional Digital Twin

The digital twin is a versioned state model, not a decorative dashboard. It connects:

- entities, ownership, governance, and authority;
- markets, customers, procurement, and revenue;
- people, employment routes, and capability creation;

- capital, restrictions, timing, cost, and dilution;
- public support, eligible cost objects, awards, and clawbacks;
- tax, transfer pricing, treaties, permanent establishments, and reporting;
- infrastructure, power, compute, land, permits, vendors, and milestones;
- evidence, obligations, incidents, transactions, outcomes, and risk;
- scenario alternatives, switch conditions, rollback, and Chronicle state.

When one object changes - ownership, a funding award, a customer contract, a new employee, an AI use, a tax rule, a power price, a deadline - the twin identifies affected claims, routes, transactions, and scenarios.

46. Evidence Docket

The Evidence Docket records the complete denominator:

- sources, retrieval dates, snapshots, facts, assumptions, and contradictions;
- models, tools, versions, prompts, environments, and deterministic logic;
- jobs, budgets, retries, failures, exceptions, and escalation;
- human and professional interventions;
- applications, approvals, signatures, filings, payments, awards, permits, and acceptance;
- complete cost, tax, delay, dilution, audit, clawback, and adverse outcomes;
- customer value, delayed effects, limitations, and externalities;
- rights, exclusions, confidentiality, publication, reuse, expiry, and revocation.

A polished narrative is not a Docket. The Docket must permit an independent reviewer to reconstruct what was known, assumed, done, spent, approved, and realized.

47. Chronicle and Selective Memory

Completed output does not automatically become institutional memory. Candidate capability k may enter Chronicle only if:

$$Admit(k) = 1,$$

where admission requires attributable evidence, rights, narrow scope, version, currentness, known failure conditions, independent challenge, reviewer acceptance, expiry, revocation, and a defined future-use boundary.

Chronicle may admit, admit with limits, repair, reject, expire, or revoke. Customer secrets, privileged material, regulated data, and third-party rights remain excluded unless an explicit protected agreement and reuse permission exist.

48. Recursive Improvement

Let g_1 be a predecessor and g_2 a successor using admitted capability. On fresh mission m_2 , define

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

where Q is quality, V verified value, T time efficiency, C cost efficiency, R risk, and H human or reviewer burden. Promotion requires a preregistered 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.

49. Trajectory Security

Persistent agents can discover methods their designers did not anticipate. The security problem spans trajectories that may include credential discovery, fragmented-secret reconstruction, sandbox escape, unauthorized communication, policy circumvention, reward hacking, hidden action sequences, or technically successful but institutionally inadmissible outcomes.

Each material scenario should be mapped as:

Threat → Precondition → Signal → Preventive Control
→ Containment → Evidence → Recovery → Owner.

Controls include unique identities, no shared root credentials, short-lived tokens, deny-by-default permissions, tool and network allowlists, spend limits, sandboxed environments, policy mediation before communications or money movement, independent telemetry, step-up authorization, automatic revocation, replay, incident response, and external red-team review.

NIST's AI RMF and Generative AI Profile provide a voluntary structure for governing, mapping, measuring, and managing AI risk, while Singapore's agentic-AI framework emphasizes reliable deployment and human accountability (National Institute of Standards and Technology, 2023; National Institute of Standards and Technology, 2024; Infocomm Media Development Authority of Singapore, 2026a). GJFFI Ω extends these ideas into fiscal, economic, evidence, and transaction authority.

Global Architecture and Illustrative Mission

The purpose of intelligence is not to admire the world. It is to allocate functions, resources, authority, and irreversible commitments into a coherent operating institution.

50. Priority Jurisdiction Archetypes

The paper does not freeze one permanent ranking. It identifies archetypes that recur in global allocation:

Table 4: Illustrative jurisdiction archetypes and activation logic.

Archetype	Illustrative nodes	Institutional role and evidence burden
Institutional and R&D root	Canada, Quebec, selected European and Asia-Pacific research nodes	Talent, R&D support, universities, compute, public legitimacy, governance, capability creation; requires ownership, control, substance, project, commercialization, and funding-interaction evidence.
Customer and capital access	United States federal and leading states, United Kingdom, Singapore, major EU markets	Enterprise sales, procurement, investors, banks, partners, acquisitions; requires nexus, data, sector, hiring, tax, securities, procurement, export, and entity-trigger analysis.
European market and programme route	EU/EEA, France, Germany, Ireland, Netherlands, Nordics	Single-market access, EIC and collaborative programmes, talent, industrial policy; requires provider/deployer roles, GDPR, state-aid, consortium, VAT, payroll, IP, and member-state overlays.
Sovereign and infrastructure capital	UAE, Saudi Arabia, Singapore, selected development-finance and sovereign nodes	Strategic capital, procurement, compute, energy, free zones, infrastructure, regional access; requires local substance, public counterparties, sanctions, foreign-investment, security, control, and project-finance evidence.
Advanced industrial ecosystem	Japan, South Korea, Germany, Switzerland, selected US and Asia nodes	Semiconductors, robotics, manufacturing, suppliers, quality, research, patents; requires localization, industrial partners, export controls, standards, procurement, workforce, and incentive terms.
Talent and scale route	India, Brazil, selected Central and Eastern European, Latin American, African, and Asia-Pacific nodes	Technical and commercial talent, large markets, cost and scale; requires payroll, classification, data, local establishment, payments, tax, currency, customer, and execution evidence.
Controlled-entry route	Mainland China, Hong Kong, other sensitive or highly regulated nodes	Market, manufacturing, capital, or supply access under special constraints; requires explicit board purpose, local counsel, data and cyber, export, sanctions, investment, IP, public-sector, and exit controls.
Functional node	Global, project SPV, standards foundation, assurance affiliate	Cross-border protocols, project finance, assurance, standards, consortium, or special-purpose functions; requires a precise legal home, governance, rights, funding, authority, and operational substance.

The correct question is not whether a node is “good.” It is whether it is good for a specified function under current facts and whether the complete architecture remains governable.

51. Illustrative Mission: AI Infrastructure Company

Consider the objective:

Discover globally. Prove locally. Allocate only after evidence passes.

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.

The Mission Constitution specifies:

- Canadian institutional root and preservation of material Canadian and Quebec eligibility where economically rational;
- US and EU customer access without premature entity proliferation;
- eighteen to thirty-six months of capital and compute planning;
- no unawarded support in base runway;
- no capability transfer without business purpose, rights, valuation, transfer-pricing, funding, and board receipts;
- no infrastructure commitment without customer, power, permit, finance, insurance, and stop gates;
- no sensitive customer or model data in the public reference environment;
- high reversibility until decisive customer and financing evidence exists.

52. Illustrative Scenario Frontier

Table 5: Illustrative scenario frontier.

Scenario	Design logic	Principal Proof Debt
Sovereign Core	Canadian and Quebec root; foreign sales without premature subsidiaries; R&D, capability, governance, and early capital concentrated.	US procurement and nexus; EU data and AI roles; facility economics; cross-border hiring.
Capital-Efficient Expansion	Core retained; US subsidiary only after customer or financing threshold; EOR for early foreign hires; customer cash and refundable support prioritized.	Canadian control; investor rights; US entity trigger; EOR-to-entity threshold; programme cash timing.
Market-Access First	US commercial entity activated early; EU partner or reseller route; R&D remains in Quebec.	Tax and IP agreements; management substance; funding interaction; data and sector obligations.
R&D and Capability First	Research and capability centered in Quebec; Canadian and European collaborative routes explored; commercial entities delayed.	Partner IP; export and data; commercialization nexus; grant covenants; customer conversion.
Infrastructure First	Project SPV selected through power, land, permit, customer, supplier, public support, and project-finance auction.	Customer commitments; permits; power; construction; financing; insurance; decommissioning.
Minimum Regulatory Surface	Remote cross-border sales, narrow data, limited external agency, minimal entities, no regulated functions.	Customer procurement acceptance; tax nexus; employment trigger; scalability and service limits.
Minimum Dilution	Customer prepayment, credits, grants, procurement, guarantees, and debt precede equity.	Award timing; working capital; customer concentration; lender covenants; founder execution capacity.

Scenario	Design logic	Principal Proof Debt
Global Functional Optimum	Every function assigned to its highest-scoring current node.	Entity proliferation, coordination, management, PE, transfer pricing, control, cost, rights, and unwind.

The system should not automatically choose the highest raw score. It should recommend the least irreversible scenario capable of proving the next consequential uncertainty while preserving extraordinary upside.

53. Capital and Public-Support Compilation

The illustrative mission may consider:

- paid pilot and customer prepayment;
- Canadian federal and Quebec R&D credits and programmes subject to exact eligibility and stacking;
- compute support or credits where current and contractually available;
- collaborative or EIC routes for European partners, where eligible;
- strategic customer or supplier capital;
- venture equity after paid proof;
- receivables, contract-backed, equipment, or infrastructure debt after cash flows exist;
- project SPV and project finance after customer, permit, power, vendor, insurance, and lender conditions.

Each resource remains in its correct evidence state. A route identified by the system is not an award. A passed screen is not administrator approval. An investor conversation is not committed capital. A term sheet is not collected cash.

54. Illustrative Cost-Object Allocation

Table 6: Illustrative cost-object architecture.

Cost object	Primary candidate route	Alternative route	Decision-changing interaction
Quebec R&D payroll	Refundable R&D credit and eligible provincial support	Federal programme or customer-funded R&D	Assistance reduces eligible base; employee, location, project, and documentation facts control.
Canadian compute	Customer-funded compute, provider credits, current public compute programme	Dedicated equipment or infrastructure facility	Programme call may be closed; data, ownership, contribution agreement, and cash timing matter.
EU demonstration	EIC or consortium route through eligible European node	Customer-sponsored pilot or distributor	Entity, consortium, partner IP, state aid, location, and commercialization terms.
US commercialization	Customer cash and procurement	Strategic capital or local incentive	Nexus, public procurement, export, entity and employment triggers.
Pilot facility equipment	Vendor finance, equipment debt, public infrastructure support	Equity or project finance	Customer commitment, permit, power, collateral, milestone, insurance, and claw-back.
Professional work	Customer cash or eligible project administration	Grant-covered professional cost where allowed	Eligibility, capitalization, related-party, procurement, tax, and audit evidence.

55. Illustrative Transaction Packet

Suppose the scenario proposes transferring software capability from the operating parent to a dedicated capability company. GJFFI Ω should not present “IP HoldCo” as a magic optimization. It should open a transaction packet containing:

- business purpose and operating rationale;
- current and future functions, assets, risks, people, and decision-making substance;
- legal ownership, chain of title, employee and contractor assignments;
- funding restrictions, customer rights, open-source, data, and export conditions;
- valuation and transfer-pricing method;
- tax, withholding, indirect tax, minimum-tax, and accounting issues;
- board approval, agreements, filings, payments, and implementation steps;
- rollback, monitoring, Evidence Docket, and Chronicle treatment.

If the transaction does not improve verified value after complete cost, risk, and unwind, the system should recommend retaining the existing structure.

56. Infrastructure and Productive Capacity

The ultimate reason jurisdiction and capital intelligence matter is not paper optimization. It is the governed conversion of resources into productive capacity:

Intelligence → Work → Proof → Authority → Capital
→ Compute, Energy, Laboratories, Industry, and Distribution.

Customer commitments should precede large irreversible assets. Project SPVs should isolate infrastructure risk. Milestone finance should condition drawdown. Permits, power, data, insurance, supply chain, cybersecurity, community, and decommissioning should exist in the evidence graph.

The Productive Horizon

The long-run opportunity is not the production of more reports. It is a trusted institutional control layer through which validated intelligence can be converted into recurring value, capital, and progressively greater productive capacity - while retaining exact evidence, authority, rights, and reversibility boundaries.

Institutional Allocation Casebook

A general allocation institution earns credibility by showing how the same proof-governed machinery changes across companies, research programmes, infrastructure, public institutions, and capital vehicles.

57. Casebook Method

The following cases are illustrative decision architectures, not recommendations to a named entity. Each begins with a Mission Constitution, generates a scenario frontier, identifies decision-changing proof, and preserves a fail-closed transaction boundary. Values are directional planning objects unless stated otherwise.

For each case, the casebook reports:

- 1. the consequential objective;
- 2. the coupled functions;
- 3. the candidate frontier;
- 4. the principal Proof Debt;
- 5. the first reversible actions;
- 6. the transaction and monitoring gates.

58. Case A: Frontier AI Company

Objective

Launch an enterprise AI company with core R&D in Montreal, customers in the United States and Europe, specialist hiring in three countries, significant compute demand, preserved founder control, and a twelve-to-eighteen-month capital plan.

58.1. Coupled decisions

The case couples Canadian and Quebec R&D support, domestic compute access, US procurement and customer contracting, European AI and data roles, founder-control status, capability ownership, transfer pricing, and institutional financing. Canada’s AI Compute Access Fund illustrates the importance of state-specific eligibility: the programme was established with up to C\$300 million, its first call closed, and 44 projects representing C\$66 million of support were announced in May 2026 (Innovation, Science and Economic Development Canada, 2026a; Innovation, Science and Economic Development Canada, 2026d). A route may therefore remain strategically relevant while current application value is zero.

Table 7: Illustrative frontier for a frontier AI company.

Architecture	Design logic	Principal Proof Debt
Montreal sovereign core	Canadian parent, Quebec R&D, foreign sales before premature subsidiaries	Customer procurement, foreign nexus, compute economics, data roles
US market first	Early US commercial entity, Canadian R&D and capability retained	Investor control, IP licence, US state tax, management substance
EU research consortium	Canadian root with EU partners and programme routes	Consortium rights, Horizon eligibility, export and data transfer, commercialization nexus
Global functional optimum	Separate best nodes for sales, R&D, hiring, capital, and compute	Entity proliferation, permanent establishment, transfer pricing, coordination and unwind

58.2. First proof actions

The first actions should test customer willingness to pay, compute availability and price, the ownership consequences of the next financing, and whether a current public-support route is actually open. The architecture should not migrate capability or create multiple subsidiaries before those uncertainties change the decision.

59. Case B: Advanced-Technology Research Consortium

Objective

Create a multinational research consortium linking universities, startups, industrial partners, and public funders while preserving publication, capability rights, commercialization options, and defensible cost allocation.

The EIC 2026 work programme allocates more than EUR 1.4 billion across Pathfinder, Transition, Accelerator, STEP Scale Up, and other schemes with materially different stage, eligibility, grant, and investment logic (European Innovation Council, 2026c; European Innovation Council, 2026a; European Innovation Council, 2026b). GJFFI Ω must therefore model the consortium as an evidence and rights graph, not a list of calls.

Table 8: Research-consortium allocation objects.

Object	Required decision	Evidence
Lead beneficiary	Which entity bears reporting, cash, and consortium liability?	Programme text, legal form, financial capacity, administrator confirmation
Work packages	Where does each experimental activity occur?	Talent, facilities, eligible expenditure, deliverables, audit trail
Background capability	What pre-existing rights enter the project?	Chain of title, licences, exclusions, confidentiality
Foreground capability	Who owns, uses, commercializes, and publishes results?	Consortium agreement, funding conditions, field-of-use rights
Cost allocation	Which costs belong to which beneficiary and programme?	Time records, invoices, allocation methodology, stacking controls
Commercialization	Which entity may create a product or spinout?	Option terms, valuations, conflicts, public-support conditions

Consortium Law

A funding award is not the objective. The objective is a research and commercialization institution whose scientific, financial, rights, and governance architecture remains coherent before, during, and after the award.

60. Case C: AI Compute and Energy Campus

Objective

Develop a large AI compute campus that combines customer demand, land, power, grid connection, equipment, cooling, fibre, public support, private capital, permits, and staged commissioning.

AI infrastructure has become a major destination of concentrated greenfield investment. UNCTAD reports that strategic sectors represented 44% of global greenfield project values in 2025 and that AI-related digital infrastructure contributed to a concentrated recovery (United Nations Conference on Trade and Development, 2026b). Canada has separately launched compute-infrastructure initiatives, including an ongoing AI Compute Challenge and a sovereign supercomputing programme (Innovation, Science and Economic Development Canada, 2026b; Innovation, Science and Economic Development Canada, 2026c).

60.1. Project architecture



The project should be ring-fenced in a special-purpose vehicle when appropriate, but only after a business-purpose case, land and power evidence, customer commitments, financing logic, environmental and community obligations, and exit or decommissioning rules exist.

Table 9: Compute-campus stage gates.

Gate	Required proof	Stop condition
Demand	Contracted or credible offtake, workload profile, service economics	Speculative capacity without customer evidence
Power	Capacity, timing, price, carbon and resilience profile	Unfunded interconnection or schedule incompatible with customers
Site and permits	Title, zoning, water, environmental, community and construction path	Fatal permitting, land, water, or community constraint
Finance	Equity, debt, guarantees, incentives, covenants, reserves	Capital stack depends on unawarded support or unrealistic utilization
Construction	EPC scope, vendors, supply chain, insurance, milestones	Unbounded cost or schedule exposure
Commissioning	Acceptance tests, security, reliability, energy and service metrics	Facility cannot meet contractual service or compliance requirements

61. Case D: Global Scientific Foundation

Objective

Create a foundation that can receive philanthropic, public, and research funding; operate programmes in multiple countries; commission scientific work; publish responsibly; and preserve mission-aligned governance.

The architecture must separate charitable or nonprofit purpose, donor restrictions, grant conditions, operating entities, local employment, programme delivery, data, publication, and capability rights. A low-tax or prestigious domicile is irrelevant if the institution cannot receive funds, employ people, conduct the programme, satisfy donor requirements, or maintain public legitimacy.

- **Root:** governance, mission lock, fiduciary duties, donor eligibility;
- **Programme nodes:** local delivery, employment, research, permits, data;
- **Funding graph:** restricted versus unrestricted funds, cost share, reporting, cash timing;
- **Rights graph:** publication, licences, open science, confidential research, community rights;
- **Evidence graph:** outputs, impact, adverse effects, independent evaluation;
- **Continuation gate:** whether the programme remains effective and mission-aligned.

62. Case E: Cross-Border Investment Fund

Objective

Launch a deep-technology investment vehicle for institutional and qualified investors across several regions while preserving a defensible management, regulatory, tax, service-provider, and portfolio-monitoring architecture.

GJFFI Ω should not treat fund domicile as a tax-ranking exercise. It must model sponsor substance, investor eligibility, marketing restrictions, management location, carried interest, governance, administration, custody, audit, tax reporting, sanctions, beneficial ownership, portfolio restrictions, and exit.

Table 10: Investment-fund architecture questions.

Function	Institutional question
Management	Where are investment decisions actually made and controlled?
Fund vehicle	Which legal form and domicile fit investors, strategy, tax, regulation, and service providers?
Capital raising	Who may be solicited, through which channels, under which exemptions or approvals?
Operations	Which administrator, auditor, bank, custodian, counsel, and compliance providers are required?
Portfolio	Which ownership, export, sanctions, sector, and foreign-investment constraints apply?
Value realization	How do distributions, withholding, carry, liquidity, and exits operate?
Monitoring	Which regulatory, tax, investor, valuation, and portfolio facts must remain current?

63. Case F: Regional Economic-Development Institution

Objective

Design a regional institution capable of attracting high-value companies, laboratories, compute, capital, and talent while preserving fiscal discipline, transparency, additionality, and public legitimacy.

The OECD’s 2026 subsidy evidence shows why public-support architecture requires a complete denominator: support can alter competition and must be assessed across grants, tax concessions, and below-market finance rather than headline announcements (Organisation for Economic Co-operation and Development, 2026c). A regional institution should optimize verified productive capacity, not awards or press releases.

The operating graph may include:

Target Sector → Project Pipeline → Incentive Offer → Private Capital
Jobs → Capability → Tax Base → Reinvestment

Every incentive offer should carry an additionality case, counterfactual, cost per verified outcome, clawback and performance terms, environmental and labour effects, conflict controls, publication rules, and post-award monitoring.

64. Cross-Case Principles

The cases differ, but the same institutional laws persist:

- 1. optimize the complete objective, not a local metric;
- 2. preserve alternative architectures;
- 3. separate discovery from eligibility, approval, collection, and audit;
- 4. use reversible actions to learn before irreversible commitments;
- 5. require professional and organizational authority for reserved judgments;
- 6. record costs, failures, contradictions, and delayed effects;
- 7. admit only scoped capability that passed;

Capital-to-Capacity Compounding

The ultimate purpose of allocation intelligence is not a prettier recommendation. It is the governed conversion of global opportunity into durable productive capacity: companies, laboratories, compute, energy, infrastructure, knowledge, and institutions capable of solving progressively harder missions.

65. From Fiscal Optimization to Productive Capacity

Tax, grants, procurement, debt, equity, and project finance are not terminal outputs. They are resource channels. Their value depends on whether they produce useful capability after full costs, restrictions, delays, and risks are measured.

Let verified contribution over period t be:

$$VC_t = \text{CashCollected}_t + \text{CostAvoided}_t + \text{RiskRetired}_t \\ + \text{OptionValueCreated}_t - \text{TotalDeliveryCost}_t - \text{ExternalityReserve}_t.$$

Only realized and rights-available contribution may be considered for reinvestment. The productive-capacity state evolves as:


$$K_{t+1} = (1 - \delta)K_t + I_t^{\text{verified}} + \text{Learning}_t - \text{Obsolescence}_t - \text{GovernanceLoss}_t,$$

where K_t includes technical, organizational, physical, financial, scientific, and institutional capability.

Reinvestment Law


Booked revenue, headline subsidy, speculative valuation, unsigned capital, and uncollected tax benefits are not free productive capital. Reinvestment follows realized contribution, applicable restrictions, liquidity, and authorized allocation.

66. The Compounding Chain



The chain compounds only if each transition remains evidence-bearing. A larger model or larger budget does not substitute for accepted proof, rights, transaction authority, or operational integration.

Capital-to-Capacity Compounding



Every transition remains evidence-bearing and revocable.

Figure 12: Capital-to-capacity compounding remains proof-governed at every transition.

67. Productive-Capacity Classes

GJFFI Ω can allocate verified resources into several capacity classes:

Table 11: Productive-capacity classes and evidence.

Capacity class	Examples	Admission evidence
Cognitive and model	Models, agents, evaluation, datasets, scientific workflows	Benchmark, safety, rights, cost, transfer, failure and monitoring evidence
Human and organiza-tional	Specialists, operators, reviewers, gov-ernance teams, provider networks	Hiring, authority, performance, conflicts, continuity and capacity evidence
Commercial	Customers, procurement access, chan-nels, licences, partnerships	Contracts, revenue, retention, acceptance, concentration and margin evidence
Financial	Cash, credit, guarantees, debt capacity, equity capacity, project finance	Collected funds, restrictions, covenants, liq-uidity, dilution and downside evidence
Physical and digital in-frastructure	Compute, energy, data centres, labs, factories, logistics, communications	Demand, permits, financing, commission-ing, resilience, environmental and service evidence

Capacity class	Examples	Admission evidence
Institutional	Entities, standards, assurance, Evidence Dockets, Chronicle capability, public-private structures	Purpose, authority, rights, auditability, recurrence and revocation evidence

68. Capacity Frontier

For capacity portfolio k , define:

$$\begin{aligned}
 PCV(k) = & \text{OutputValue}(k) + \text{StrategicOption}(k) + \text{Learning}(k) + \text{Resilience}(k) \\
 & - \text{CapitalCost}(k) - \text{OperatingCost}(k) - \text{Delay}(k) \\
 & - \text{Externality}(k) - \text{LockIn}(k) - \text{Unwind}(k).
 \end{aligned}$$

The system selects a capacity frontier rather than one asset. A small amount of compute, customer-funded validation, specialist hiring, and rights-cleared data may outperform a premature large facility. Conversely, once demand, power, permits, and finance are proven, a project-specific infrastructure vehicle may create more durable capacity than continued cloud rental.

69. Public Support as Capacity Catalyst

Public support should be evaluated through productive additionality:

$$\text{Additionality} = \text{Outcome}_{\text{with support}} - \text{Outcome}_{\text{credible counterfactual}}.$$

The relevant questions are:

- Does the support cause a project, accelerate it, enlarge it, de-risk it, or merely subsidize what would occur anyway?
- Which private capital is mobilized?
- Which capabilities, jobs, infrastructure, tax base, scientific outputs, or public benefits result?
- What location, IP, employment, environmental, reporting, or clawback obligations attach?
- Is the project still economically coherent if support is delayed, reduced, or denied?

The OECD's industrial-subsidy work reinforces the need to treat grants, tax concessions, and below-market finance as one competitive architecture rather than separate headline instruments (Organisation for Economic Co-operation and Development, 2026c).

70. Compute, Energy, and Sovereign Infrastructure

AI systems convert electricity, chips, networks, data, engineering, and capital into useful computation. Allocation must therefore connect software economics to physical constraints. Canada, the European Union, Singapore, the United States, and other jurisdictions increasingly combine compute policy, public finance, domestic capacity, and governance (Innovation, Science and Economic Development Canada, 2026b; Innovation, Science and Economic Development Canada, 2026c; European Innovation Council, 2026c; Infocomm Media Development Authority of Singapore, 2026b).

A compute allocation should model:

- workload profile and elasticity;
- hardware availability and depreciation;
- energy, cooling, water, grid, and resilience;

- data location, security, export, and customer restrictions;
- cloud, colocation, owned, sovereign, and hybrid options;
- grant, credit, debt, guarantee, and project-finance interactions;
- customer commitments, utilization, service levels, and obsolescence;
- decommissioning and environmental obligations.

71. Capital Allocation Council

Capital allocation is a governed decision plane, not an automatic consequence of model output. The council receives:

- verified contribution and liquidity;
- mission priorities and capacity gaps;
- candidate investments and opportunity cost;
- restrictions, covenants, and funding conditions;
- expected return, tail risk, externality, and reversibility;
- evidence quality and responsible-owner recommendation.

The council may allocate to product, people, data, compute, infrastructure, distribution, assurance, or reserves. It may also decline to reinvest if the evidence does not support productive use.

Capital-to-Capacity Objective

The institution seeks to expand the frontier of useful, rights-controlled, resilient productive capacity. The measure of success is not the amount of capital attracted, but the verified capability and public or customer value created per unit of constrained capital, time, energy, and institutional attention.

72. Network and Civilization-Scale Externalities

At sufficient scale, allocation decisions affect energy systems, labour markets, scientific ecosystems, housing, communities, supply chains, security, and public finances. The objective function must therefore include externalities and distribution, not merely private enterprise value.

For affected stakeholder class h , let $W_h(s)$ be welfare impact and B_h a minimum acceptable boundary. An admissible large-scale architecture should satisfy:

$$W_h(s) \geq B_h \quad \forall h \in \mathcal{H}_{material},$$

subject to transparent trade-offs and responsible public authority where public power is involved.

Potential positive externalities include scientific capacity, resilient infrastructure, clean energy, regional productivity, skilled employment, public services, technology diffusion, and new industrial capability. Potential harms include concentration, displacement, resource consumption, unequal access, surveillance, security risk, fiscal dependency, and stranded assets.

73. Compounding Without Overreach

The long-run ambition is legitimate only if the proof discipline scales with it. Three safeguards are decisive:

1. **No automatic authority expansion.** More capital or more capable models do not enlarge permission.
2. **No automatic memory expansion.** A successful result does not become universal doctrine.
3. **No automatic reinvestment.** Realized value remains subject to restrictions, reserves, governance, and competing needs.

Discover globally. Prove locally. Allocate only after evidence passes.

Productive Intelligence Law. Verified intelligence should cause bounded work; accepted evidence may support authorized action; authorized action may create protected capability and recurring value. Only realized, rights-available value may finance the next capacity frontier.

Governance, Legal Boundary, and Institutional Protection

The system is strongest when it minimizes unnecessary regulatory surface through actual architecture while remaining ready to satisfy mandatory law when real functions, data, people, capital, and transactions require it.

74. Authority Boundary

GJFFI Ω supplies research, source retrieval, comparison, simulation, orchestration, evidence, completeness testing, monitoring, and transaction control. It does not convert a model output into legal, tax, fiduciary, regulatory, investment, immigration, audit, engineering, certification, procurement, or government authority.

The default public-reference posture is:

- browser-local decision support;
- no custody or movement of money;
- no signing, filing, contract execution, or autonomous external communication;
- no professional-client relationship by implication;
- no high-impact automated decision about a natural person;
- no public-power exercise;
- no confidential, privileged, personal, regulated, classified, credential, vulnerability, or unauthorized third-party data;
- no implication that private notices bind regulators or waive mandatory law;
- candidate recommendations separated from organizational authority;
- binding action conditioned on current sources, exact facts, responsible review, approval, and receipts.

Actual-Function Law

Actual users, functions, data, rights, ownership, control, communications, counterparties, jurisdictions, external effects, and money flows determine legal treatment. Labels, disclaimers, or entity names do not transform regulated conduct into unregulated conduct.

75. Fifteen Legal and Institutional Control Planes

Table 12: Integrated control planes.

Control plane	Core questions
AI classification and roles	What system, use, sector, provider, deployer, importer, distributor, product, public-authority, or high-impact role exists?
Authority and permissioning	Who may observe, access, alter, communicate, deploy, purchase, recommend, approve, disclose, or escalate?
Privacy and data	What personal, sensitive, confidential, training, customer, worker, or cross-border data exists, under what lawful basis and role?
Cybersecurity and resilience	What identities, credentials, tools, supply chain, logging, incident, recovery, and continuity controls apply?
Capability, IP, copyright, provenance	Who created and owns capability, data, code, models, output, inventions, rights, licences, and reusable methods?
Consumer protection and competition	Are claims, pricing, ranking, referrals, terms, dark patterns, market power, or representations unfair or misleading?
Employment and workplace	Does the system recruit, monitor, evaluate, classify, pay, dismiss, or affect workers or contractors?
Financial services and settlement	Does the activity involve custody, payment, credit, investment, insurance, securities, digital assets, advice, or fiduciary duty?

Control plane	Core questions
Health and life sciences	Does the system affect diagnosis, treatment, clinical research, devices, patient data, safety, or regulated claims?
Government, defense, and procurement	Are public funds, tenders, classified information, national security, local content, sovereign data, or public powers involved?
Sanctions, export, and investment review	Are parties, owners, technologies, destinations, end uses, licences, foreign investments, or controlled goods implicated?
Tax and transfer pricing	What tax, treaty, permanent establishment, withholding, indirect, payroll, minimum-tax, valuation, and intercompany issues arise?
Language, accessibility, localization	What official-language, consumer, accessibility, documentation, cultural, records, and local-presence requirements apply?
Records, audit, and evidence	What must be logged, retained, signed, reproduced, reported, independently reviewed, or made available to authorities?
Incident and regulator engagement	Who detects, contains, investigates, reports, notifies, remediates, preserves evidence, and communicates with affected parties or authorities?

76. AI and Agentic Governance

The EU AI Act creates role- and risk-specific obligations through a staged timeline that continued to evolve in 2026. Article 50 transparency obligations apply from 2 August 2026, with current guidance distinguishing provider and deployer duties; the 2026 political agreement places key standalone high-risk rules in December 2027 and product-embedded high-risk rules in August 2028 (European Commission, 2026a; European Commission, 2026c; European Commission, 2026b; European Commission, 2026d). GJFFI Ω should classify whether a deployment acts as provider, deployer, importer, distributor, product manufacturer, employer, public authority, or regulated-sector actor; and whether the use is prohibited, high-risk, transparency-regulated, general-purpose, or ordinary decision support.

NIST’s AI RMF emphasizes governance, context mapping, measurement, and management, while the Generative AI Profile adds risks and actions for generative systems (National Institute of Standards and Technology, 2023; National Institute of Standards and Technology, 2024). NIST’s 2026 AI Agent Standards Initiative and critical-infrastructure profile work emphasize secure interoperability, measurement, and operational trust, while Singapore’s updated agentic-AI framework adds concrete controls for bounding powers, meaningful accountability, multi-agent systems, and third-party agents (National Institute of Standards and Technology, 2026a; National Institute of Standards and Technology, 2026b; Infocomm Media Development Authority of Singapore, 2026b). These sources support a common principle: capable agents remain subordinate to accountable organizational controls.

A production GoalOS deployment should maintain:

- system, model, tool, provider, subprocessor, and data registers;
- risk and impact assessments appropriate to role and use;
- objective, authority, data, tool, spend, and external-action bounds;
- evaluation, red-team, incident, change, and rollback records;
- human review, override, escalation, and responsible ownership;
- transparency, explanation, records, retention, and affected-person controls;
- professional and regulator cooperation where applicable.

77. Privacy and Data Architecture

The public reference edition should collect no confidential or sensitive mission facts. Protected deployments require a data map, roles, lawful basis, minimization, purpose limitation, security, processor and subprocessor review, transfer mechanism, localization, retention, deletion, rights response, incident response, and customer agreement.

A no-secrets gate should detect and reject:

- credentials, authentication tokens, API keys, secrets, and private keys;
- privileged legal or tax communications;
- personal, biometric, health, financial, children, worker, or other regulated data;
- controlled technology, classified, export-restricted, or national-security information;
- vulnerabilities, exploit details, or unauthorized third-party information;
- customer data for which the user lacks disclosure and processing authority.

78. Sanctions, Export Controls, and Investment Screening

A global allocation system can increase exposure by proposing more jurisdictions, partners, capital sources, and technology flows. Each route should screen:

- parties, beneficial owners, banks, intermediaries, and counterparties;
- country, region, sector, technology, item, software, model, data, end use, and end user;
- ownership and control, blocked property, facilitation, diversion, and evasion risk;
- export classification, licence exceptions, deemed exports, services, cloud, and technical assistance;
- foreign direct investment, national-security, critical-infrastructure, and sector approvals;
- contractual representations, audit, termination, reporting, and ongoing monitoring.

The system should never optimize around sanctions or export restrictions through hidden intermediaries, false end users, artificial routing, fragmented credentials, or disguised ownership.

79. Professional Authority

GoalOS may identify issues, retrieve sources, compare options, prepare dossiers, coordinate providers, test completeness, record decisions, and monitor expiry. Licensed and responsible providers supply legal opinions, tax opinions, audits, valuations, regulated investment activity, immigration representation, engineering sign-off, official certifications, and reserved judgments.

Every active jurisdiction should have named professional roles rather than a generic “review required” label. A professional receipt records:

- professional identity, firm, qualification, jurisdiction, role, and conflict;
- exact question, facts, sources, assumptions, scope, limitations, and conclusion;
- effective date, expiry, material-change triggers, responsible signatory, and reliance boundary;
- linked transaction, scenario, evidence objects, and unresolved issues.

80. Conflicts, Marketplace, and Ranking Independence

Provider compensation, referral fees, marketplace economics, strategic investment, sponsorship, or procurement commissions must not silently influence jurisdiction or scenario ranking. The architecture should separate:

Independent Scenario Engine → Recommendation → Disclosed Provider Marketplace.

A conflict receipt should identify relationships, compensation, alternatives, ranking independence, user selection, and review. Where separation is impossible, the conflict becomes visible Proof Debt and can disqualify the recommendation.

81. Claims and Publication Controls

The institution should distinguish:

- discovered route;
- sourced route;
- preliminarily eligible route;
- verified eligibility;
- professional conclusion;
- authorized application or transaction;
- approved award or financing;
- collected cash or completed activation;
- audited or monitored realized value.

No public claim should imply an award, approval, compliance, certification, partnership, customer result, tax saving, funding amount, or superiority that the Evidence Docket cannot support.

No-Immunity Law

Terms, notices, licences, and architecture can allocate rights, establish conditions, document boundaries, and reduce avoidable exposure. They cannot bind regulators or non-signatories, waive mandatory law, guarantee immunity, or change the legal effect of actual conduct.

82. Deployment Modes

GJFFI Ω should support escalating modes:

1. **Protected Reference Mode:** browser-local, non-confidential research and simulation; no external authority.
2. **Controlled Internal Mode:** organization-approved users, data controls, logs, internal authority, no regulated external action.
3. **Professional-Assisted Mode:** protected customer facts, scoped providers, professional receipts, contractual controls.
4. **Regulated Planning Mode:** regulated issues mapped and execution disabled until applicable licensed and organizational authority exists.
5. **Separate Execution Environment:** external agents, filings, money, custody, sensitive data, public power, or critical infrastructure governed through a distinct reviewed system.


Production Implementation Architecture

The paper becomes an operating institution only when sources, claims, scenarios, authority, transactions, outcomes, and revocation are represented in a live, secure, versioned system.

83. Production System Topology

A production GJFFI Ω deployment should separate at least nine services:

1. **Mission service:** Mission Constitutions, versions, permissions, and acceptance criteria;
2. **Evidence graph:** facts, claims, sources, contradictions, expiry, and lineage;
3. **Registry service:** jurisdiction nodes, instruments, professionals, providers, and control planes;
4. **Eligibility and fiscal engine:** rule evaluation, stacking, tax-base interaction, cash timing, and audit logic;
5. **Scenario engine:** architecture generation, frontier preservation, simulation, regret, and switch conditions;
6. **Agentic work fabric:** bounded AGI Jobs, tools, environments, budgets, traces, and verifiers;
7. **Transaction controller:** receipts, approvals, holds, signatures, implementation state, and rollback;
8. **Chronicle:** selective capability admission, version, rights, expiry, and revocation;
9. **Monitoring and notification:** source change, deadline, covenant, entity, funding, tax, infrastructure, and evidence alerts.



84. Event-Sourced Institutional State

Material changes should be append-only events rather than silent overwrites. Examples include:

- source retrieved, changed, superseded, or expired;
- claim created, supported, contradicted, accepted, or revoked;
- instrument state changed;
- mission version frozen;
- scenario generated or challenged;
- receipt approved, rejected, or expired;
- transaction blocked, authorized, implemented, or unwound;
- capability admitted, narrowed, repaired, or revoked;
- realized outcome recorded.

The current state is a projection of the event log. This supports audit, rollback, reproducibility, and historical comparison.

85. The Evidence Graph

A claim object should contain:

Table 13: Minimum claim-object fields.

Field	Meaning
Claim	Exact proposition that may affect a decision
Scope	Mission, entity, project, jurisdiction, function, period, and transaction
Source set	Official, professional, customer, market, technical, and internal evidence
Provenance	Retrieval, author, date, version, transformations, and access rights
Contradictions	Evidence inconsistent with the claim
Status	Candidate, sourced, verified, professionally reviewed, authorized, realized, revoked
Materiality	Decision impact if wrong
Confidence and residual uncertainty	Structured uncertainty without converting it into permission
Owner and reviewer	Accountable human or institution
Expiry and refresh rule	Date, trigger, or event requiring revalidation
Downstream dependencies	Scenarios, instruments, jobs, transactions, or capabilities affected

86. Source Ingestion and Change Detection

The monitoring layer should:

1. retrieve official pages, gazettes, regulator notices, programme guides, procurement portals, and administrator communications;
2. preserve a timestamped snapshot and cryptographic digest;
3. extract candidate changes while retaining the original source;
4. classify the change as opening, deadline, amount, eligibility, authority, tax, law, guidance, enforcement, status, or clarification;
5. identify affected claims, instruments, scenarios, transactions, and customers;
6. route material changes to an accountable reviewer;
7. supersede or revoke prior evidence only through a versioned decision;
8. rerun affected scenarios after approval.

Freshness Policy

Static dates are insufficient. Each source receives a refresh class: event-driven, daily, weekly, monthly, quarterly, annual, or transaction-triggered. A high-impact programme deadline may require daily monitoring; a constitutional treaty fact may require annual or event-triggered review.

87. Policy-as-Code and Human Authority

Rules should be machine-readable where possible but never imply that the rule engine replaces interpretation. A rule object contains:

Condition → Result → Evidence → Reviewer → Expiry.

Examples:

- entity must be incorporated in the jurisdiction before application;

- foreign ownership above a threshold requires additional review;
- eligible expenditure must be incurred after a stated date;
- assistance received reduces the tax-credit base;
- an IP transfer requires valuation and transfer-pricing receipts;
- a public-sector or critical-infrastructure use triggers a legal control plane.

Rules may output pass, fail, unresolved, or professional review required. They do not output legal immunity.

88. Agentic Work Fabric

Each AGI Job is isolated by mission, data, tools, workspace, budget, and output schema. A production job record includes:

- objective and material claim;
- authorized inputs and forbidden data;
- source class and retrieval method;
- model, prompt, tools, code, and environment versions;
- time, spend, and compute budget;
- expected artifact and validation schema;
- independent verifier context;
- stop, escalation, and recovery route;
- Evidence Docket return path;
- rights and retention rules.

NIST's 2026 AI Agent Standards Initiative emphasizes secure interoperability for systems that act on users' behalf, while Singapore's updated agentic-AI framework emphasizes bounding powers, meaningful human accountability, lifecycle controls, and end-user responsibility (National Institute of Standards and Technology, 2026a; Infocomm Media Development Authority of Singapore, 2026b). GJFFI Ω applies these principles to consequential fiscal and institutional work.

89. Trajectory Security

Security must evaluate the full multi-step trajectory, not only individual prompts. Controls include:

- unique agent and tool identities;
- short-lived credentials and least privilege;
- allowlisted tools, domains, data stores, communications, and transactions;
- sandboxed code, browser, and file execution;
- independent telemetry outside the agent's control;
- credential scanning and reconstruction detection;
- spend, transaction, and rate limits;
- human step-up approval for external or irreversible action;
- automatic revocation and containment;
- replayable traces and incident evidence.

Adversarial Design Assumption

A capable system may discover a technically effective path that the institution did not intend. Success therefore cannot be defined solely by task completion. The path, permissions, evidence, and institutional acceptability must also pass.

90. Professional Network and Responsibility Map

Every active jurisdiction should have named coverage for relevant disciplines:

- corporate and commercial law;
- tax and transfer pricing;
- employment and immigration;
- public funding and procurement;
- accounting and audit;
- sanctions, export controls, and investment review;
- privacy, cybersecurity, and AI regulation;
- infrastructure, energy, engineering, permitting, and insurance.

The network record should disclose jurisdiction, discipline, licence or qualification, conflicts, commercial relationship, scope, rates, service levels, currentness, and performance evidence. Provider compensation must remain structurally separated from scenario ranking.

91. Deployment Modes

Table 14: Deployment modes and authority boundaries.

Mode	Permitted use	Boundary
Public reference	Non-confidential research, simulation, education, architecture exploration	No customer secrets, filing, transaction, regulated advice, or external autonomous action
Protected enterprise	Contracted data, role-based access, customer-specific evidence, professional collaboration	Requires security, privacy, retention, subprocessors, incident and authority controls
Regulated workflow	Reserved activities coordinated with licensed or authorized providers	GoalOS does not impersonate reserved authority; execution follows provider and legal controls
Sovereign / critical infrastructure	Public or strategic missions, classified or sensitive environments, infrastructure and public authority	Separate environment, enhanced assurance, procurement, security, public-law and political-accountability architecture

92. Operational Cadence

A production institution should operate at several cadences:

Continuous: source monitoring, deadline alerts, ownership and eligibility changes, transaction holds, security and incident detection;

Weekly: new opportunities, material legal or programme changes, Proof Debt, source expiry, mission exceptions;

Monthly: funding and capital committee, cost-object stacking, application pipeline, realized-value review;

Quarterly: jurisdiction auction, scenario rerun, entity topology, capital sequence, provider performance, infrastructure options;

Annual or trigger-based: substance, transfer pricing, IP, insurance, employment, sanctions, security, governance, full architecture reset.

93. A 1,000-Day Production Path

Table 15: Illustrative production path.

Period	Institutional objective	Required evidence
Days 0–30	Constitute product, data, authority, claims, source and professional boundaries	Versioned Mission Contract, source spine, role map, first customer pipeline, legal and security baseline
Days 31–100	Complete three comparable paid missions	Paid scope, acceptance criteria, evidence docket, external review, implementation plan
Days 101–250	Repeat and standardize	Comparable economics, reusable schemas, decline rules, provider coverage, rights-cleared capability
Days 251–500	Add live monitoring and transaction workflows	Source changes, alerts, professional receipts, application and transaction state, repeat purchase
Days 501–750	Portfolio and infrastructure scale	Multi-entity state, provider network, project SPV workflows, capital and procurement integration
Days 751–1,000	Earn category authority	Independent benchmark, realized customer value, recurring offices, audited outcomes, fresh successor improvement

94. Institutional KPIs

The operating dashboard should prioritize:

- verified customer value realized;
- support identified, applied, approved, collected, and retained after audit;
- dilution avoided or financing quality improved;
- decision time and implementation time;
- incomplete transactions blocked;
- source freshness and evidence expiry;
- Proof Debt retired and remaining;
- professional-review cycle time;
- false eligibility and false ineligibility rates;
- Chronicle admission, repair, revocation, and fresh-transfer performance;
- trust incidents and material claim corrections;
- renewal and Expansion Office conversion.

Production Readiness Law

A polished interface is not production readiness. Production requires live evidence, exact facts, bounded execution, independent challenge, responsible authority, transaction integration, realized outcomes, revocation, and operational ownership.

Discover globally. Prove locally. Allocate only after evidence passes.

Commercialization and Institutional Scale

The economic prize is not a one-time report. It is the recurring control layer through which institutions allocate functions, discover resources, authorize transactions, and adapt to global change.

95. Product Architecture

GJFFI Ω can support a layered product system:

1. **Readiness Scan:** identify complexity, material opportunity, immediate stop conditions, and evidence gaps.
2. **Global Architecture Mission:** produce the jurisdiction allocation, scenario frontier, capital and support architecture, Proof Debt, and 100-day plan.
3. **Funding and Fiscal Mission:** build cost-object, eligibility, stacking, tax-interaction, application, and cash-timing architecture.
4. **Entity and Transaction Mission:** compare structures and prepare complete professional and board evidence packets.
5. **Infrastructure and Procurement Mission:** coordinate site, compute, energy, supplier, public support, finance, and milestones.
6. **Global Expansion Office:** monitor source drift, funding, tax, entity substance, permanent establishment, transfer pricing, hiring, regulation, capital, procurement, and infrastructure.
7. **Portfolio and Sovereign Editions:** govern multiple entities, projects, portfolio companies, sectors, or regional development missions.
8. **Protected Intelligence API:** expose rights-controlled, evidence-dated objects and workflows to approved systems.

96. Illustrative Product Ladder

Table 16: Illustrative commercial architecture; not a current offer, forecast, or guarantee.

Product	Illustrative price	Principal economic value
Global Allocation Readiness Scan	Free–US\$10K	Qualify mission complexity, value at stake, and stop conditions.
Jurisdiction and Architecture Mission	US\$25K–150K	Function allocation, scenario frontier, evidence and authority map.
Funding and Fiscal Mission	US\$75K–500K	Eligibility, stacking, tax interaction, capital sequence, application architecture.
Entity and Transaction Mission	US\$100K–500K	Structure comparison, professional coordination, evidence packets, transaction gating.
Infrastructure and Public-Finance Mission	US\$250K–2M+	Site, energy, compute, permits, support, procurement, SPV, and financeability.
Global Expansion Office	US\$250K–1M+ yearly	Monitoring, successor frontiers, Proof Debt, and transaction control.
Portfolio Edition	US\$500K–3M+ yearly	Repeated missions across a fund, group, accelerator, university, or platform.
Sovereign or Strategic Programme	US\$5M–20M+	Regional or national investment, infrastructure, incentive, and institutional allocation.
Protected Intelligence API	Usage-based	Evidence-dated source, eligibility, scoring, monitoring, and workflow infrastructure.

These price ranges are management hypotheses. The appropriate fee depends on value at stake, complexity, number of jurisdictions and functions, professional burden, infrastructure, risk, evidence depth, and execution scope.

97. Customer Value Gate

A disciplined qualification rule is

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

where F is the customer fee. This does not guarantee return. It prevents scarce institutional capacity from being assigned to missions whose plausible value cannot justify the cost.

Value may arise from support, credits, financing, procurement, reduced dilution, lower cost of capital, faster activation, avoided entity or tax failure, improved customer or investor confidence, lower professional burden, better option value, and reusable rights-cleared capability.

98. The Recurring Global Expansion Office

A jurisdictional decision is episodic; the architecture drifts continuously. The recurring product monitors:

- funding openings, deadlines, awards, amendments, and clawbacks;
- tax, treaty, minimum-tax, incentive, and reporting changes;
- entity substance, management, beneficial ownership, and governance;
- permanent establishment, transfer pricing, withholding, and intercompany flows;
- employment, payroll, immigration, contractor, and EOR triggers;
- investor, cap-table, financing, covenant, and foreign-control events;
- AI, data, sector, sanctions, export, procurement, and investment-screening change;
- infrastructure, power, compute, supplier, permit, and public-finance opportunities;
- source expiry, transaction authority, and unresolved Proof Debt.

The renewal case is evidence-based: what changed, which assumptions became stale, which opportunities opened, which risks were retired, which transactions were blocked, which routes were realized, and what the successor frontier now recommends.

99. Business Potential

The opportunity spans professional advisory, tax and legal research, public funding, formation, global employment, private-capital administration, procurement, site selection, and infrastructure. The distinctive position is not replacing every provider. It is deciding which combination should exist, in what sequence, under which evidence and authority, and with what continuing control.

Table 17: Illustrative operating stages; not forecasts.

Stage	Operating model	Illustrative annual revenue	Proof required
Elite institution	Productized high-value missions plus early recurring offices	US\$8M–30M	Paid outcomes, repeatability, professional review, realized value.
Platformized institution	Standardized missions, live monitoring, portfolio customers, workflow integration	US\$75M–300M	Recurring revenue, outcome data, provider coverage, strong gross margin.
Global category leader	Hundreds of recurring offices, missions, APIs, and network economics	US\$500M–1.5B	Durable retention, verified benchmarks, trusted transaction control.
Institutional standard	Thousands of offices, strategic programmes, infrastructure, professional network, global evidence graph	Multi-billion scale	Category standard, deep integrations, capital-allocation competence, exceptional trust.

The largest outcome depends on becoming the recurring control plane, not on selling static reports.

100. Illustrative Scale Configurations

A decacorn operating configuration could include:


- 500 Global Expansion Offices at an average US\$1 million;
- US\$250 million of architecture, fiscal, transaction, and infrastructure missions;
- US\$150 million of protected API and workflow usage;
- US\$100 million of disclosed marketplace, validation, and partner economics.

This produces an illustrative US\$1 billion annual scale. It is not a valuation forecast.

A megacorn-candidate configuration could include thousands of recurring offices, major strategic and infrastructure programmes, global professional coverage, protected data and workflow networks, and billions in annual revenue. The scale thesis is structurally possible only if GJFFI Ω demonstrates superior realized allocation, high recurrence, strong margins, low trust failure, and a defensible evidence graph.


101. The Institutional Moat

The moat is not the list of countries, programmes, or rates. The durable asset is the rights-cleared outcome graph:



Over time, this graph can reveal which programmes award which profiles; how long routes take; which combinations survive audit; which professional concerns change decisions; which jurisdictions introduce hidden cost; which entity patterns create burden; which infrastructure assumptions fail; which early signals predict rejection or clawback; and which capabilities transfer successfully.

The compounding sequence is:



102. The First Commercial Wedge

The strongest initial segment is an institutionally financed AI, deep-tech, clean-tech, life-sciences, industrial-technology, or infrastructure company making coupled decisions about root and foreign entities; R&D and capability location; hiring; public support and compute; capital for eighteen to thirty-six months; US and European market entry; and infrastructure or project-finance preparation.

The initial offer can be framed as a *Global AI Company Allocation Mission*: a paid, bounded mission producing the defensible architecture, support and capital stack, transaction map, Evidence Docket, and 100-day command sequence.

103. A 1,000-Day Institutional Path

Table 18: Illustrative 1,000-day proof-gated path.

Phase	Days	Exit condition
Constitute	0–30	Product, authority boundary, data model, professional baseline, and first paid pipeline.
Prove	31–100	Three accepted paid missions, independent review, and one expansion purchase.
Repeat	101–250	Standardized segment playbooks, recurring monitoring, and measurable mission economics.
Execute	251–500	Live source operations, provider network, realized funding or fiscal routes, transaction packets.
Network	501–750	Portfolio customers, API, benchmarks, rights-cleared capability, and partner ecosystem.
Institution	751–1,000	Diversified recurring cash flows, audited controls, validated benchmarks, and capital-market readiness.

Scale Law

Scale evidence, currentness, recurrence, and realized customer value before scaling payroll, entity count, infrastructure, or public claims.

GJFFIBench Ω and the Proof Programme

A state-of-the-art claim is earned only when the institution outperforms appropriate alternatives on paid, fresh, consequential missions after the complete denominator is measured.

104. GJFFIBench Ω

GJFFIBench Ω is a matched-mission benchmark for global allocation. It compares:

1. a sponsor-led point-solution stack;
2. a multidisciplinary expert team;
3. a static frontier-model report;
4. an ungated multi-agent workflow;
5. proof-governed GJFFI Ω .

The benchmark should include mission classes such as:

- AI and deep-tech company allocation;
- R&D and capability location;
- public-support and fiscal stacking;
- foreign market entry and hiring;
- infrastructure and compute site selection;
- capital and project-finance architecture;
- nonprofit or research consortium design;
- economic-development and sovereign investment allocation.

105. Benchmark Metrics

Table 19: GJFFIBench evaluation dimensions.

Metric	Evaluation
Decision quality	Independent expert and sponsor assessment of architecture, alternatives, assumptions, and decision relevance.
Verified value	Conservative expected and subsequently realized customer value after complete costs, taxes, delay, dilution, audit, clawback, and adverse outcomes.
Evidence completeness	Material claims with current primary sources, exact facts, contradictions, professional receipts, and expiry.
Opportunity recall	Material programmes, routes, procurement, capital, talent, and infrastructure alternatives found.
Precision	Proportion of routes that survive exact eligibility and professional review.
Time	Elapsed time to admissible decision, application, activation, and realized outcome.
Cost	Models, data, labour, professional, application, travel, filing, compliance, financing, and coordination cost.
Reviewer burden	Hours, complexity, and cognitive load required from sponsor, board, and specialists.
Transaction safety	Incomplete, contradictory, unauthorized, stale, or economically inferior transactions prevented.
Tail performance	P10, CVaR, stress, concentration, and hostile-world outcome.
Reversibility	Stranded cost, exit time, termination, data and people portability, and unwind exposure.

Metric	Evaluation
Chronicle precision	False admission, false rejection, scope error, stale reuse, rights failure, and revocation performance.
Successor transfer	Improvement on fresh missions after controlling for budget, information, and evaluator conditions.
Institutional regret	Value of the best feasible alternative discovered after the decision, including omitted scenarios.
Trust failure	Material legal, fiscal, security, data, funding, claim, transaction, or professional-authority breach.

106. Experimental Protocol

A credible benchmark should:

1. preregister mission facts, hidden facts, constraints, value, budget, tools, time, and acceptance criteria;
2. select independent evaluators before execution and disclose conflicts;
3. normalize or disclose differences in source and professional access;
4. preserve outputs, logs, costs, retries, failures, interventions, and complete denominators;
5. evaluate both immediate recommendation and post-implementation outcome;
6. publish negative results, omitted alternatives, and governance failures;
7. repeat on fresh mission classes and changing conditions;
8. prevent the production team from selecting only favourable cases;
9. compare the successor against the predecessor under controlled conditions;
10. distinguish designed, verified, authorized, completed, realized, and audited claims.

107. Commercial Proof Run

The first external proof should remain narrow.

Table 20: Minimum evidence gates for the first external proof programme.

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


A public-safe Receipt Pack should include paid scope, preregistered acceptance, mission facts, comparator, evidence summary, independent verdict, complete cost denominator, implementation result, realized value, adverse outcomes, limitations, repeat purchase, Chronicle decision, and exact claim.

108. Research Agenda

The architecture suggests an empirical programme:

1. **Scenario quality:** compare scenario frontiers with independent multidisciplinary teams.
2. **Value of Proof:** estimate which evidence most changes allocation decisions.
3. **Eligibility precision:** compare preliminary screening with administrator and professional outcomes.
4. **Funding realization:** compare discovered gross support with applied, awarded, collected, audited, and retained net cash.
5. **Stacking accuracy:** test cost allocation, tax-base, assistance, minimum-tax, and clawback interaction.
6. **Graph efficiency:** measure when parallel AGI Jobs reduce elapsed time and when coordination dominates.
7. **Verifier independence:** test blind evaluation, context separation, evaluator diversity, and conflict controls.
8. **Transaction safety:** measure prevented incomplete, contradictory, stale, or unauthorized actions.
9. **Chronicle precision:** estimate false admission, false rejection, stale reuse, and revocation rates.
10. **Successor transfer:** determine whether admitted capability improves fresh missions without governance regression.
11. **Recurring value:** measure whether the Expansion Office reduces risk, cost, delay, missed opportunity, and regret.
12. **Externalities:** measure labour, environmental, infrastructure, concentration, regulatory, and public-value effects.

109. Claim Ladder



A route may be called designed when generated, sourced when the official authority is identified, verified when current terms and exact facts support it, reviewed when applicable qualified authority has examined it, authorized when the responsible organization approves, completed when the action occurs, realized when the external value or risk reduction occurs, and monitored while obligations and currentness remain under control.

State-of-the-Art Claim Law

No complete denominator, no state-of-the-art claim. Superior architecture must be demonstrated against appropriate alternatives after cost, risk, authority, security, currentness, and realized outcomes are measured.

Limitations and Falsification Boundary

Ambition becomes credible when the institution states where it can fail, what remains unproven, and how it will know.

110. Limitations

110.1. Coverage is not completeness

A 70-node registry and 463 governed routes do not imply universal or continuously current coverage. Subnational rules, sector overlays, unpublished administrator practice, language, local facts, programme amendments, political change, and new instruments can alter the result. The system must expose missing coverage and source age rather than convert scale into certainty.

110.2. Official routes are not open programmes

The 420 official-authority routes are governed discovery obligations. They identify where current procurement, tax, grant, compute, export, and financing opportunities should be sought. They are not automatically open calls, confirmed eligibility, or available cash.

110.3. Scores are not truth

Node and scenario scores depend on the mission, weights, normalization, evidence, and proxies. They structure comparison and direct evidence acquisition. They cannot replace exact legal, fiscal, commercial, engineering, or financing analysis.

110.4. Probabilities are difficult to calibrate

Award, customer, permitting, financing, tax, construction, and geopolitical probabilities are uncertain. Tail events and regime shifts can invalidate historical experience. Robust optimization reduces fragility but does not eliminate Knightian uncertainty (Knight, 1921; Ben-Tal et al., 2009).

110.5. Professional judgment is not fungible

Legal, tax, fiduciary, audit, immigration, public-authority, engineering, and financing judgments depend on facts, jurisdiction, responsibility, and professional standards. A receipt increases accountability; it does not make every conclusion correct or comparable.

110.6. A browser-local release is not a live intelligence service

The public reference application is a deterministic, browser-local planning institution over a dated snapshot. It does not browse or refresh sources in real time, submit applications, form entities, sign contracts, move money, or execute external action. These limitations reduce attack surface and increase inspectability, but they prevent the standalone from being a continuously operated production service.

110.7. Recursive systems can amplify error

A permissive Chronicle can turn one successful case into universal doctrine. An excessively restrictive Chronicle can prevent useful transfer. Admission precision, rights, expiry, revocation, and fresh-successor testing remain empirical problems.

110.8. The benchmark may be gamed

Any published metric can become a target (Goodhart, 1975; Campbell, 1979). GJFFIBench requires evaluator independence, hidden cases, multiple mission classes, complete denominators, negative-result publication, and periodic redesign.

110.9. Commercial scale is not a present capability claim

The paper includes high-scale operating configurations because the category may support substantial recurring economics. Their inclusion does not establish product-market fit, revenue, valuation, awards, capital raised, customer outcomes, or global standard status.

111. Principal Risks and Controls

Table 21: Principal risks, failure modes, and controls.

Risk	Failure mode	Control
False legal or tax certainty	Strategic output is treated as binding authority.	Current sources, exact facts, professional receipt, expiry, and transaction gate.
Grant theatre	Unawarded support enters runway or drives the project.	Discovered, eligible, applied, approved, committed, collected, and audited states separated.
Official-route inflation	Discovery routes are represented as live named programmes.	Registry-class disclosure, route-level evidence discount, programme verification before reliance.
Shell-company optimization	Tax score dominates substance and operating reality.	Business purpose, people and risk map, entity-cost penalty, local substance and transfer-pricing review.
Scenario monoculture	One model anchors on one attractive structure.	Diverse archetypes, Pareto frontier, independent red team, counterfactuals, minimax regret, switch conditions.
Metric gaming	System maximizes tax, grant size, or number of entities.	Verified Global Allocation Value includes coordination, delay, audit, risk, externality, and unwind.
Professional conflict	Adviser or marketplace compensation biases ranking.	Compensation disclosure, separated ranking and referral functions, independent review, conflict receipt.
Data leakage	Sensitive facts enter public or third-party systems.	No-secrets gate, protected deployment, minimization, access control, retention, deletion, incident response.
Capability overreach	One successful case becomes universal memory.	Narrow scope, rights, version, known failures, expiry, revocation, fresh transfer.
Stale-source drift	An old law, programme, or deadline remains embedded.	Refresh queue, evidence decay, blocked reliance, affected-scenario rerun, Chronicle supersession.
Agentic trajectory escape	Agent discovers unauthorized credentials, tools, communications, or actions.	Deny-by-default identity, allowlists, short-lived tokens, sandboxing, telemetry, step-up authorization, revocation.
Infrastructure overinvestment	Assets are funded before demand, permits, power, or finance.	Customer commitments, project SPV, milestones, permits, lender conditions, stop and unwind gates.
Funding clawback	Support is received but conditions later fail.	Covenant register, evidence calendar, restricted cash, audit reserve, ongoing monitoring.
Tax incentive leakage	Assistance reduces another credit or triggers top-up tax.	Cost-object model, tax-base adjustment, minimum-tax and safe-harbour analysis, net value.
Founder or expert dependence	Every exception returns to one individual.	Fixed products, schemas, delegated authorities, provider network, explicit escalation and decline paths.

Risk	Failure mode	Control
Regulatory evasion	Architecture disguises actual regulated conduct.	Non-evasion constitution, actual-function classification, professional authority, blocked execution.

112. Conclusion

The world already contains mature tools for tax research, public funding, company formation, employment, capital administration, procurement, site selection, cloud, infrastructure, and professional advice. GJFFI Ω does not need to replace them all. It can coordinate them around the outcome the institution actually wants.

Tell GoalOS what the institution must accomplish.
It turns the world into an evidence-bearing resource graph - and the objective into a governed global allocation.

The paper’s strongest claim is institutional rather than technological. Better models make research, drafting, screening, simulation, and orchestration cheaper. The durable opportunity is to convert that capability into a governed system of objectives, facts, sources, instruments, eligibility, interactions, scenarios, evidence, authority, selective memory, recurring control, and productive reinvestment.

Final Institutional Law

No consequential objective, no mission. No exact facts, no eligibility. No current official source, no material reliance. No cost object, no stacking result. No complete denominator, no net-value claim. No responsible authority, no transaction. No Chronicle admission, no future influence. No fresh comparison, no improvement claim. No realized outcome, no commercial authority.

Schemas and Operating Doctrine

The institution becomes executable only when objectives, facts, routes, receipts, rights, and review clocks are represented as governed objects.

A. Fifteen-Factor Jurisdiction Constitution

Table 22: Minimum jurisdiction-node constitution.

Field	Required content
Identity	Stable node ID; jurisdiction; subnational or special-zone scope; currency; languages; time zone; record version.
Strategic role	Candidate uses: institutional root, sales, hiring, R&D, capability ownership, capital, public support, infrastructure, or value realization.
Customer access	Demand, procurement access, contracting friction, payments, localization, channels, and support economics.
Talent	Depth, cost, immigration, classification, payroll, benefits, retention, universities, laboratories, and specialist ecosystems.
R&D incentives	Refundability, eligible expenditure, ownership restrictions, nexus, collaboration, cash timing, audit, and clawback.
Commercialization incentives	Scale-up support, procurement, export, demonstration, manufacturing, customer, and infrastructure incentives.
Effective tax	Corporate, withholding, indirect, payroll, property, top-up, incentives, losses, timing, and uncertainty.
Capability and IP	Assignment, ownership, DEMPE, valuation, licensing, funding, open source, export, and data controls.
Financing	Customer prepayment, grants, debt, strategic capital, venture and growth equity, development finance, guarantees, and project finance.
Procurement	Public and private procurement access, qualification, local content, security, language, insurance, and past performance.
AI and data regulation	Roles, privacy, transfers, cybersecurity, records, sector rules, automated decisions, transparency, and incidents.
Export controls	Sanctions, parties, investment screening, controlled technology, end use, end users, licences, and diversion risk.
Energy and infrastructure	Power, compute, land, connectivity, water, logistics, permits, resilience, suppliers, and commissioning.
Cost and coordination	Formation, annual operation, professionals, reporting, banking, payroll, travel, translation, management, and entity proliferation.
Reversibility	Exit, liquidation, migration, termination, employee transition, grant unwind, IP and data portability, and stranded cost.
Treaties and stability	Tax and investment treaties, enforceability, institutions, political and legal stability, currency, and capital controls.
Evidence state	Sources, retrieval dates, confidence, contradictions, professional owner, expiry, next review, and activation status.

B. Governed Instrument Record

Table 23: Minimum instrument schema.

Field	Required content
Instrument identity	Stable ID, official name, jurisdiction, authority, type, official source, route class, and status.
Opening and timing	Opening, deadline, cadence, decision, reimbursement, expiry, and last verified.
Budget and award	Programme envelope, minimum and maximum award, cost share, match, guarantee, interest, or investment range.
Eligible entity	Residence, establishment, ownership, control, size, age, revenue, headcount, sector, consortium, public/private status.
Eligible project	Technology readiness, project location, activities, milestones, expenditure, start date, and environmental or social requirements.
Capability and IP	Ownership, licence, access, publication, commercialization, localization, foreground/background IP, and change of control.
Procurement route	Competitive call, challenge, framework, direct award, tax filing, automatic entitlement, lender, or investor process.
Interaction and stacking	Assistance caps, cost allocation, federal/provincial/state interaction, basis reduction, other aid, safe harbours, double dip.
Tax and accounting	Taxability, credit, basis, income recognition, capitalization, currency, withholding, and financial-statement treatment.
Clawback and covenants	Location, headcount, spend, ownership, reporting, audit, retention, security, scope change, repayment, and default.
Cash and execution	Application owner, evidence checklist, runway, reimbursement risk, application cost, approval uncertainty, and next action.
Evidence state	Discovered, sourced, screened, verified, reviewed, applied, approved, committed, collected, audited, expired, rejected, or revoked.

C. Eligibility Rule Set

A route should, where applicable, test:

1. residence, establishment, registration, and project location;
2. ownership, beneficial ownership, foreign control, and change-of-control conditions;
3. employee count, revenue, balance sheet, age, stage, and group aggregation;
4. sector, prohibited activity, controlled technology, and regulated-industry status;
5. technology readiness, novelty, collaboration, and commercialization pathway;
6. eligible expenditure, incurred date, payroll, compute, equipment, subcontractors, travel, and overhead;
7. matching capital, cost share, project size, liquidity, and runway;
8. capability ownership, publication, licensing, data, open-source, and localization conditions;
9. consortium, academic, customer, supplier, regional, or public-sector partner requirements;
10. prior assistance, cumulation, basis interaction, intensity, and double-dip restrictions;
11. deadline, lead time, project start, approval timing, and reimbursement timing;
12. sanctions, export, debarment, security, beneficial ownership, evidence, and professional review.

D. Cost-Object Stacking Ledger

Table 24: Minimum stacking record.

Object	Required content
Cost object	ID, entity, project, location, vendor or employee, date, currency, accounting class, evidence, and total amount.
Proposed route	Instrument, allocated amount, eligible percentage, cost share, application status, and expected cash date.
Alternative route	Competing route, eligibility, expected value, timing, burden, and reason not selected.
Interaction rule	Assistance cap, prohibited overlap, exclusion, ordering, tax-base or credit reduction, and top-up effect.
Net value	Gross support less interaction, tax leakage, professional and compliance cost, delay, audit reserve, and clawback reserve.
Decision	Chosen route, authority, rationale, evidence state, review date, stop conditions, and Chronicle reference.
Audit trail	Invoices, payroll, time records, contracts, payment, allocation method, applications, awards, correspondence, and amendments.

E. Transaction Receipt Set

A material transaction requires, where applicable:

1. business-purpose memorandum;
2. legal memorandum;
3. tax memorandum;
4. valuation;
5. transfer-pricing analysis;
6. funding-interaction analysis;
7. board or delegated approval;
8. implementation checklist;
9. Chronicle receipt.

Each receipt records the exact question, facts, scope, sources, conclusion, limitations, owner, reviewer, effective date, expiry, conditions, related objects, and revocation triggers.

F. Evidence Docket Schema

Table 25: Core Evidence Docket objects.

Object	Required content
Objective record	Sponsor, decision, value at stake, deadline, acceptance, authority, and exclusions.
Fact register	Verified facts, source, date, owner, confidence, expiry, and contradiction.

Object	Required content
Assumption ledger	Assumption, sensitivity, evidence needed, decision impact, and retirement plan.
Claims matrix	Material claim, supporting evidence, contradiction, status, reviewer, and allowed use.
AGI Job ledger	Worker, tools, data, budget, output, failure, verifier, and Docket return.
Cost ledger	Models, data, professionals, applications, travel, filings, capital, tax, delay, and adverse outcomes.
Authority register	Approval, scope, owner, effective date, expiry, conditions, delegation, and revocation.
Transaction receipts	Agreements, filings, payments, awards, permits, financing, implementation, and acceptance.
Outcome record	What happened, customer value, adverse outcome, delayed effect, and residual risk.
Rights register	Customer, platform, employee, contractor, professional, open-source, third-party, and reuse rights.
Chronicle decision	Admit, admit with limits, repair, reject, revoke, expire, or supersede.
Successor result	Baseline, fresh mission, metrics, improvement, regressions, verdict, and scope.

G. Source Freshness and Verification States

Every material record carries source date, retrieval date, official status, effective date, evidence owner, confidence, contradictions, expiry, next review, and invalidation triggers. A route can progress from discovered to sourced, screened, verified, reviewed, authorized, realized, and audited; or regress to stale, repair required, rejected, expired, or revoked.

H. GJFFIBench Mission Record

A benchmark record should contain sponsor class, sector, objective, value, deadline, authority; authorized and hidden facts; baseline modes and budgets; evaluator selection and conflicts; architectures and omitted alternatives; route eligibility, Proof Debt, and professional dependencies; complete cost and time; decision and implementation; realized and adverse outcomes; institutional regret; Chronicle decision; exact public claim; and successor test.

I. Canonical Operating Doctrine

Discover globally.
Decompose function by function.
Represent every route as a governed object.
Test eligibility against exact facts.
Allocate every cost only once.
Sequence capital by evidence state.
Generate competing architectures.
Simulate consequential uncertainty.
Challenge adversarially.
Resolve locally.

Govern institutionally.
Execute only within authority.
Record the complete denominator.
Admit only what passed.
Revoke when stale.
Reinvest only realized, verified value.

Final Institutional Law

No consequential objective, no mission. No exact facts, no eligibility. No current official source, no material reliance. No cost object, no stacking result. No complete denominator, no net-value claim. No responsible authority, no transaction. No Chronicle admission, no future influence. No fresh comparison, no improvement claim. No realized outcome, no commercial authority.

J. Mission Constitution v2 Schema

Field group	Required content
Objective	Consequential objective, lawful purpose, value at stake, deadline, success and failure criteria
Sponsor and authority	Accountable sponsor, board or delegated authority, affected stakeholders, decision and signing rights
Entity facts	Current entities, residence, ownership, control, employees, revenue, locations, regulated activities, tax and reporting status
Function priorities	Relative importance of sell, hire, R&D, capability, capital, support, infrastructure, and value realization
Constraints	Prohibited jurisdictions, data and sovereignty rules, export and sanctions limits, maximum entity complexity, dilution and risk limits
Resources	Available cash, matching capital, debt capacity, team capacity, compute, facilities, professional budget, runway
Evidence policy	Source classes, freshness thresholds, reviewer types, acceptance, contradictions, uncertainty and expiry
Execution policy	Permitted tools, environments, communications, spend, transactions, counterparties and escalation
Rights	Confidentiality, customer rights, platform rights, third-party rights, publication, retention, deletion and reuse
Monitoring	Trigger events, review cadence, responsible owners, successor conditions, stop and rollback rules

K. Jurisdiction Node v2 Schema

A node should contain:

- identity, region, tier, legal level, currency, languages, and relevant subnational overlays;
- scores and evidence for all fifteen factors by function;
- intended institutional roles and prohibited or high-risk uses;
- entity, employment, tax, data, AI, IP, procurement, sanctions, export, investment, and infrastructure triggers;
- treaty and cross-border interaction notes;

- official-source spine and professional owners;
- current activation state, expiry, next review, stop conditions, and reversibility;
- realized mission outcomes and known failure patterns.

L. Instrument v2 Schema

Field	Required content
Identity	Instrument, programme, authority, jurisdiction, category, source IDs
Availability	Opening, deadline, rolling or closed status, budget, allocation mechanism
Entity eligibility	Residence, size, ownership, control, sector, age, revenue, funding stage
Project eligibility	Location, TRL, activity, commercialization path, eligible period, consortium, end use
Cost eligibility	Salaries, contractors, equipment, compute, capital, overhead, taxes, exclusions
Economics	Award or finance range, intensity, cost share, matching, cash timing, reimbursement, tax treatment
Interactions	Assistance caps, stacking, double-dip, tax-base reduction, minimum-tax, accounting and procurement effects
Strategic conditions	IP, employment, production, ownership, security, export, localization, reporting
Risk	Audit, clawback, security, guarantees, covenants, change of control, cancellation, currency
Evidence	Application dossier, professional review, administrator confirmation, award, claim, payment and audit receipts
State	Discovered, sourced, screened, verified, reviewed, applied, approved, committed, collected, audited, rejected, expired, revoked
Ownership	Application owner, reviewer, executive sponsor, deadline owner, next action

M. Proof Debt Taxonomy

Class	Examples	Retirement evidence
Fact debt	Ownership, revenue, employee, location, spend, use or deadline unknown	Verified entity records, contracts, ledgers, declarations, technical facts
Source debt	Outdated page, secondary summary, inaccessible or contradictory authority	Current official source, archived snapshot, administrator confirmation
Interpretation debt	Rule applicability or interaction unclear	Scoped professional opinion, regulator or programme clarification
Economic debt	Timing, tax leakage, cost share, dilution, audit or unwind not modelled	Complete financial and cash-flow model with sensitivities
Authority debt	Decision owner, licence, board approval, signatory or delegation unclear	Authority register, board or professional receipt
Execution debt	Provider, filing, contract, permit, bank, payroll or technical pathway unresolved	Implementation plan, provider acceptance, filing or contract evidence
Rights debt	Ownership, licence, confidentiality or reuse unclear	Assignments, licences, rights register, consent, exclusions
Outcome debt	Recommendation not implemented or value not realized	Transaction, payment, award, customer, activation, audit and outcome receipts

N. Scenario Dossier v2

Each scenario dossier should include:

1. executive thesis and mission fit;
2. function-by-function jurisdiction allocation;
3. entity, ownership, management, and intercompany topology;
4. capital sequence and public-support portfolio;
5. cost-object stacking ledger;
6. regulatory, tax, data, sanctions, export, employment, and infrastructure issue map;
7. Verified Global Allocation Value and complete denominator;
8. uncertainty, P10, CVaR, regret, concentration, reversibility, and switch conditions;
9. Proof Debt, responsible owner, evidence action, cost, and deadline;
10. first 30-, 100-, and 365-day action graph;
11. transaction-gate requirements;
12. monitoring and rollback constitution;
13. plain-language recommendation and reasons not to choose the scenario.

O. Authority Object

An authority object should record:

- person or institution;
- legal or organizational source of authority;
- scope, jurisdiction, transaction, amount, and time;
- conditions, conflicts, delegations, and prohibited actions;
- evidence required before exercise;
- effective date, expiry, suspension, and revocation;
- signature or approval receipt;
- downstream actions affected by loss of authority.

P. Source Snapshot Object

Field	Content
Source identity	Authority, title, URL, document number, jurisdiction and category
Acquisition	Retrieval date, method, agent, access condition, archive location and digest
Version	Publication date, effective date, amendment, consultation, proposal, final or repealed state
Extracted claims	Exact claims and locations within the source
Applicability	Entity, project, sector, threshold, territorial and temporal scope
Contradictions	Other official or professional sources inconsistent with the claim

Freshness	Review class, expiry date, trigger event and owner
Rights	Copyright, access, quotation, storage and distribution conditions
Dependencies	Instruments, nodes, scenarios, transactions and customers affected

Q. Capital Allocation Receipt

A capital allocation receipt should state:

1. verified contribution and liquidity available;
2. restrictions, covenants, reserves, and excluded amounts;
3. candidate capacity investments and alternatives;
4. expected value, tail risk, externality, irreversibility and opportunity cost;
5. evidence and reviewer conclusion;
6. approving authority and limits;
7. implementation milestones and stop conditions;
8. realized output and post-investment review;
9. Chronicle treatment.

R. GJFFIBench v2 Trial Record

A matched trial should record:

- mission facts and blinded benchmark arm;
- expert team, point-solution stack, static model, ungated agents, or GJFFI configuration;
- budget, elapsed time, compute, professional and reviewer cost;
- architectures generated and alternatives preserved;
- evidence completeness, contradictions, and Proof Debt;
- transaction errors prevented or introduced;
- customer decision and implementation;
- support, tax, financing, hiring, market, or infrastructure outcome;
- realized customer value and adverse effects;
- Chronicle admission and fresh successor performance;
- material trust incidents.

S. Glossary

Admissible architecture: A scenario whose current evidence and authority requirements have passed for the decision stage.

Authority edge: A graph edge that blocks progression until responsible permission exists.

Chronicle: Selective, versioned, rights-governed institutional memory.

Cost object: A defined expenditure unit that may be allocated to funding, tax, procurement, or capital routes.

Evidence Docket: The complete record of facts, sources, assumptions, jobs, costs, contradictions, approvals, outcomes, and rights.

Official-authority route: A governed discovery channel to an authority; not necessarily a currently open instrument.

Proof Debt: A material unresolved claim, fact, source, authority, interaction, execution, rights, or outcome dependency.

Realized support: Cash or equivalent value actually collected and retained subject to applicable restrictions and audit.

Scenario frontier: The set of non-dominated candidate institutional architectures.

Transaction Gate: The receipt and authority system that determines whether a material action may proceed.

Verified Global Allocation Value: Mission value after the complete cost, risk, delay, dilution, coordination, and unwind denominator.

T. Sovereign Fiscal Operating Doctrine v2

Canonical Doctrine

Discover globally. Constitute precisely. Decompose function by function. Preserve competing architectures. Price the complete denominator. Separate routes from awards and awards from cash. Allocate each cost once. Search for counterexamples. Prove locally. Require responsible authority. Execute fail-closed. Record the full truth. Admit only what passed. Revoke when stale. Reinvest only verified contribution. Generate a successor only on a fresh mission.

References

- Ben-Tal, Aharon, Laurent El Ghaoui, and Arkadi Nemirovski (2009). *Robust Optimization*. Princeton University Press.
- Campbell, Donald T. (1979). “Assessing the Impact of Planned Social Change”. In: *Evaluation and Program Planning* 2.1, pp. 67–90.
- Coase, Ronald H. (1937). “The Nature of the Firm”. In: *Economica* 4.16, pp. 386–405.
- Dixit, Avinash K. and Robert S. Pindyck (1994). *Investment under Uncertainty*. Princeton University Press.
- European Commission (2026a). *AI Act: Regulatory Framework and Implementation Timeline*. URL: <https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai> (visited on 07/24/2026).
- (2026b). *Guidelines on the Implementation of the Transparency Obligations for Certain AI Systems under Article 50 of the AI Act*. URL: <https://digital-strategy.ec.europa.eu/en/library/guidelines-transparency-obligations-providers-and-deployers-ai-systems> (visited on 07/24/2026).
- (2026c). *Simplification and Phased Application of High-Risk AI Rules*. URL: <https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai> (visited on 07/24/2026).
- (2026d). *Transparency Obligations under Article 50 of the AI Act*. URL: <https://digital-strategy.ec.europa.eu/en/faqs/transparency-obligations-under-article-50-ai-act> (visited on 07/24/2026).
- European Innovation Council (2026a). *EIC Accelerator 2026*. URL: https://eic.ec.europa.eu/eic-funding-opportunities/eic-accelerator_en (visited on 07/24/2026).
- (2026b). *EIC STEP Scale Up 2026*. URL: https://eic.ec.europa.eu/eic-funding-opportunities/step-scale_en (visited on 07/24/2026).
- (2026c). *EIC Work Programme 2026*. URL: https://eic.ec.europa.eu/eic-funding-opportunities/eic-2026-work-programme_en (visited on 07/24/2026).
- Goodhart, Charles A. E. (1975). “Problems of Monetary Management: The U.K. Experience”. In: *Papers in Monetary Economics*. Reserve Bank of Australia.

- Infocomm Media Development Authority of Singapore (2026a). *Model AI Governance Framework for Agentic AI*. URL: <https://www.imda.gov.sg/resources/press-releases-factsheets-and-speeches/factsheets/2026/updated-model-ai-governance-framework-for-agentic-ai> (visited on 07/24/2026).
- (2026b). *Updated Model AI Governance Framework for Agentic AI*. URL: <https://www.imda.gov.sg/resources/press-releases-factsheets-and-speeches/factsheets/2026/updated-model-ai-governance-framework-for-agentic-ai> (visited on 07/24/2026).
- Innovation, Science and Economic Development Canada (2026a). *AI Compute Access Fund*. URL: <https://ised-isde.canada.ca/site/ised/en/canadian-sovereign-ai-compute-strategy/ai-compute-access-fund/application-guide-ai-compute-access-fund> (visited on 07/24/2026).
- (2026b). *AI Compute Challenge*. URL: <https://ised-isde.canada.ca/site/ised/en/canadian-sovereign-ai-compute-strategy/ai-compute-challenge> (visited on 07/24/2026).
- (2026c). *Canada Launches National Initiative to Build Large-Scale AI Supercomputing Capacity*. URL: <https://www.canada.ca/en/innovation-science-economic-development/news/2026/04/canada-launches-national-initiative-to-build-large-scale-ai-supercomputing-capacity.html> (visited on 07/24/2026).
- (2026d). *Government of Canada Supports 44 Canadian Companies through the AI Compute Access Fund*. URL: <https://www.canada.ca/en/innovation-science-economic-development/news/2026/05/government-of-canada-supports-44-canadian-companies-using-ai-to-transform-industries-and-create-jobs.html> (visited on 07/24/2026).
- Knight, Frank H. (1921). *Risk, Uncertainty and Profit*. Boston: Houghton Mifflin.
- National Institute of Standards and Technology (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. Tech. rep. NIST AI 100-1. U.S. Department of Commerce. URL: <https://www.nist.gov/itl/ai-risk-management-framework> (visited on 07/24/2026).
- (2024). *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile*. Tech. rep. NIST AI 600-1. U.S. Department of Commerce. URL: <https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-generative-artificial-intelligence> (visited on 07/24/2026).
- (2026a). *AI Agent Standards Initiative for Interoperable and Secure Innovation*. URL: <https://www.nist.gov/news-events/news/2026/02/announcing-ai-agent-standards-initiative-interoperable-and-secure> (visited on 07/24/2026).
- (2026b). *Concept Note: AI RMF Profile on Trustworthy AI in Critical Infrastructure*. URL: <https://www.nist.gov/programs-projects/concept-note-ai-rmf-profile-trustworthy-ai-critical-infrastructure> (visited on 07/24/2026).
- Organisation for Economic Co-operation and Development (2026a). *Global Minimum Tax: Implementation and Administrative Toolkit*. URL: <https://www.oecd.org/en/topics/sub-issues/global-minimum-tax.html> (visited on 07/24/2026).
- (2026b). *Inclusive Framework on BEPS: Pillar Two Administrative Guidance and Side-by-Side Package*. URL: <https://www.oecd.org/en/topics/sub-issues/global-minimum-tax.html> (visited on 07/24/2026).
- (2026c). *Industrial Subsidies Reach Highest Levels since the Global Financial Crisis*. URL: <https://www.oecd.org/en/about/news/press-releases/2026/06/industrial-subsidies-reach-highest-levels-since-the-global-financial-crisis-says-oecd.html> (visited on 07/24/2026).
- (2026d). *The Global Minimum Tax Implementation Toolkit*. OECD Forum on Tax Administration. URL: https://www.oecd.org/en/publications/the-global-minimum-tax-implementation-toolkit_0ee66d84-en.html (visited on 07/24/2026).
- Rockafellar, R. Tyrrell and Stanislav Uryasev (2000). “Optimization of Conditional Value-at-Risk”. In: *Journal of Risk* 2.3, pp. 21–41.
- United Nations Conference on Trade and Development (2026a). *Global Investment Rises 6 Percent to USD 1.6 Trillion, but Development Gains Remain Uneven*. URL: <https://investmentpolicy.unctad.org/news/hub/1797/20260707-global-investment-rises-6-to-1-6-trillion-but-development-gains-remain-uneven> (visited on 07/24/2026).
- (2026b). *World Investment Report 2026 and Global Investment Trends*. URL: <https://unctad.org/publication/world-investment-report-2026> (visited on 07/24/2026).
- Williamson, Oliver E. (1985). *The Economic Institutions of Capitalism*. New York: Free Press.