MONTREAL.AI · QUEBEC.AI

Gym, Specialist ASI & Successor Ω

The Sovereign Mission Environment,
the Proven Intelligence Institution,
and the Successor to the Successor in Singularity Time

Alpha, Beta, Spread, and the Mission-Sovereign Institution under Fresh Proof

Vincent Boucher
President, MONTREAL.AI & QUEBEC.AI
https://montreal.ai/
president@montreal.ai

Version 2.1.0 · Definitive Institutional Edition · 10 August 2026

Montreal, Quebec, Canada

THE CONSTITUTIONAL OVERTURE

From rented intelligence to a proof-bearing mission institution

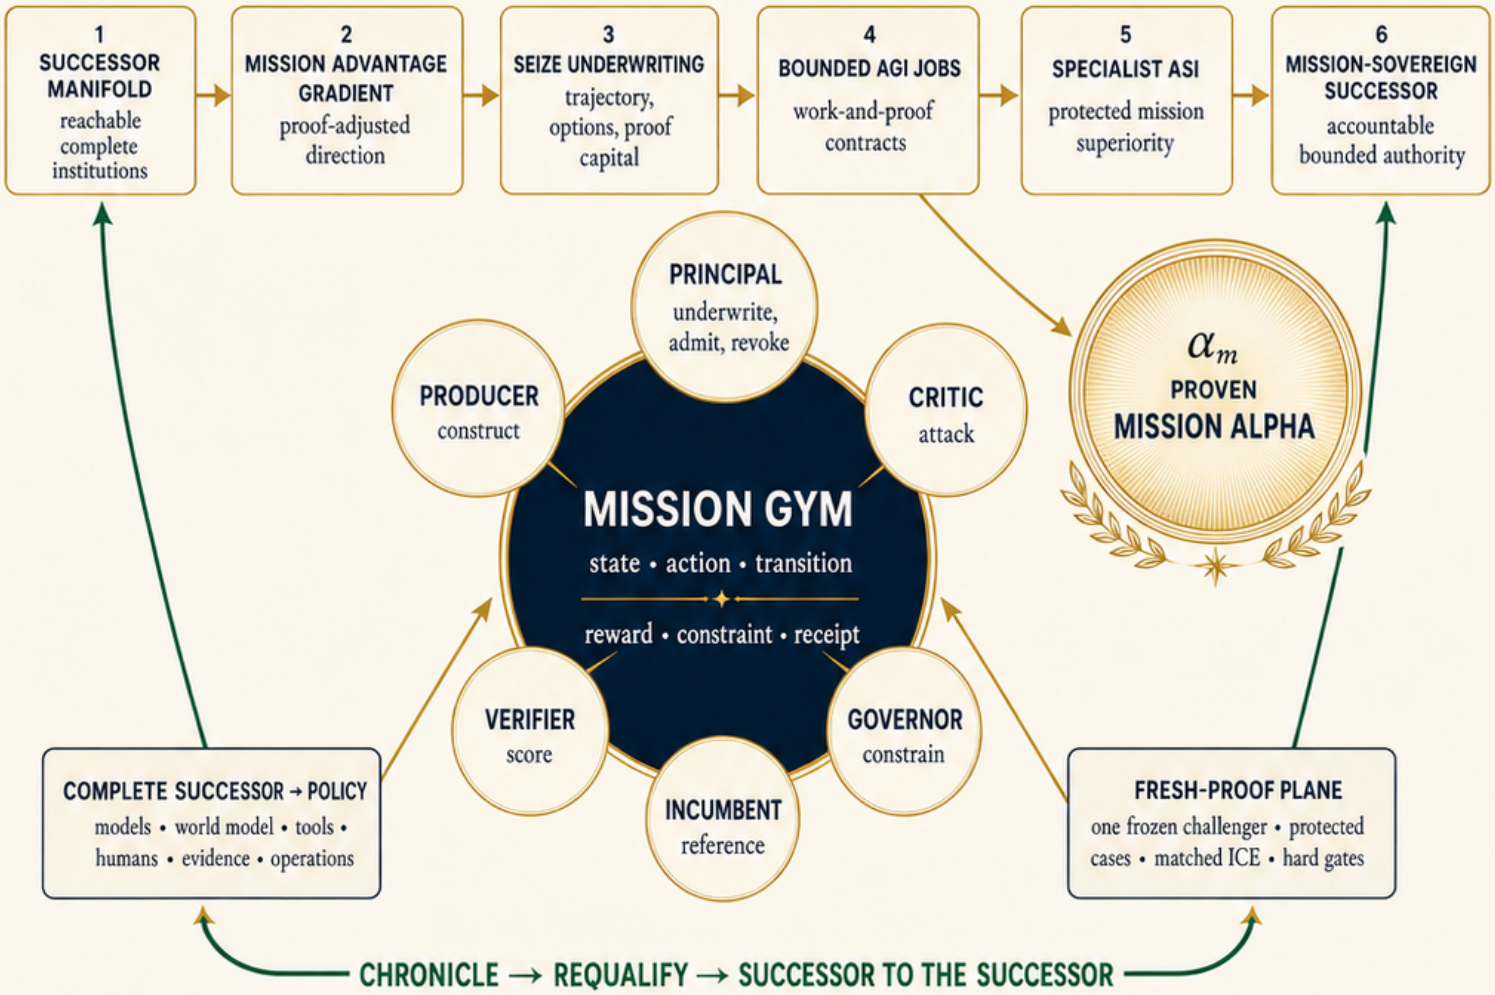

Mission Alpha

$$\alpha_m = U_m(S_m) - \max\{U_m(I_m), U_m(B_m)\} - C_m - R_m - D_m - \varepsilon_{gym}.$$

The Spread is only an opportunity. Alpha exists only after the complete denominator and protected comparison.

$$A_t \le \phi(E_t).$$

Capability may justify candidacy. Fresh proof may establish mission dominance. Only accountable admission may grant authority.

The complete doctrine

Beta is rented. The Spread is underwritten. Alpha is manufactured. Specialist ASI is proven. Authority never exceeds fresh proof. The Gym teaches and tests; reality proves; accountable authority admits; Chronicle prepares the successor to the successor.

Governing declaration

In calendar time, institutions optimize plans. In Singularity Time, institutions govern succession. The Gym is the sovereign representation of the mission; the policy is only one acting component; Specialist ASI is an earned comparative state; the Mission-Sovereign Successor is the complete admitted institution; and the deepest recursion is the verified succession of the successor itself.

Abstract

This paper defines the relationship among the Mission Gym, Specialist ASI and Mission-Sovereign Successor as the executable core of the **GoalOS** Verified Succession Institution. The Mission Gym is a controlled, versioned and rights-cleared representation of a consequential mission : its hidden state, role-specific observations, bounded actions, transition dynamics, reward vector, hard constraints, scenario distributions, verifier and proof receipts. It is not merely a reinforcement-learning benchmark. It is the mission's sovereign experimental environment : a corporate asset that can be used to form, compare, attack, prove, requalify and replace complete successor architectures.

The paper distinguishes the policy acting inside the Gym from the complete Successor. A policy may select actions, but a successor institution includes models, world models, data, retrieval, deterministic tools, planners, specialist roles, independent verifiers, human principals, workflow integration, evidence custody, operational controls, rollback, Chronicle and an Authority Envelope. A Specialist ASI is the complete mission-specific intelligence institution that outperforms the incumbent and the best credible general, specialist, human, software or hybrid alternative on fresh, protected and representative work after quality, critical error, complete economics, human burden, rights, sovereignty, resilience and governance are counted. A Mission-Sovereign Successor is that proven system only after separate accountable admission and the grant of scoped, versioned, monitored, expiring and revocable authority.

We formalize the Mission Gym as a constrained partially observable stochastic game, the Successor Manifold as a metric architecture space, the Mission Advantage Gradient as a proof-adjusted value direction, and SEIZE as a Bayesian real-options and least-action trajectory selector. The multiplayer game includes the principal, incumbent, producer, critic, verifier, governor, affected parties, providers and adversaries. The mechanism separates formation, evaluation, acceptance, admission and capital allocation so that no claimant becomes the final judge of its own authority. The paper develops environment basis risk, policy–institution separation, reward grounding, gauge-invariant mission sovereignty, information-geometric architecture distance, phase transitions, Alpha half-life, proof-collateralized authority and a recursive successor operator that generates the successor to the successor without allowing self-certification or automatic authority inheritance.

The canonical progression remains :

Successor Manifold → Mission Advantage Gradient → SEIZE Underwriting → Bounded AGI Jobs → Specialist ASI → Mission-Sovereign Successor.

The paper provides a complete Gym stack, a twenty-one-job successor-manufacturing portfolio, a multiplayer mechanism, a protected Fresh-Successor Gate, an Authority Lattice, Chronicle and Negative Capability Graphs, machine-readable records, a reference API, a worked synthetic real-estate invoice-integrity succession cycle, corporate use cases, commercial productization and falsification conditions. Its governing rule is exact : authority never exceeds fresh proof.

Keywords : GoalOS; Mission Gym; Successor Gym; Specialist ASI; Mission-Sovereign Successor; Successor Manifold; Mission Advantage Gradient; SEIZE; multiplayer game; stochastic game; recursive institutional improvement; Alpha; Beta; Spread; proof capital; Chronicle; Authority Envelope; Singularity Time.

Résumé exécutif

Ce document définit la relation entre le *Mission Gym*, l'ASI spécialisée et le successeur souverain pour la mission. Le Gym constitue une représentation contrôlée, versionnée et vérifiable de la mission : état caché, observations propres à chaque rôle, actions permises, transitions, récompenses multidimensionnelles, contraintes, scénarios adverses, vérificateur et reçus de preuve. Il devient un actif souverain de l'entreprise parce qu'il permet de substituer les modèles et les fournisseurs sans perdre la mission, les tests, l'évidence, la mémoire ni l'autorité.

L'ASI spécialisée n'est pas simplement la politique qui agit dans le Gym. Il s'agit de l'institution d'intelligence complète qui dépasse l'architecture en place et la meilleure solution crédible sur du travail frais, protégé et représentatif, après prise en compte de la qualité, des erreurs critiques, des coûts complets, de la charge humaine, des droits, de la souveraineté, de la résilience et de la gouvernance. Le successeur souverain pour la mission est cette intelligence prouvée, mais seulement après admission séparée par une autorité responsable, avec une enveloppe d'autorité circonscrite, versionnée, temporaire, surveillée et révocable.

La récursion centrale porte sur le successeur du successeur. Le système actuel peut proposer des descendants, produire de nouvelles architectures et demander une promotion ; il ne peut toutefois pas consulter les cas protégés, modifier la constitution après les résultats, se certifier, s'installer ni hériter automatiquement de l'autorité. Chronicle conserve uniquement les capacités, échecs et preuves qui survivent à la réalité. La règle directrice demeure : l'autorité ne dépasse jamais la preuve fraîche.

Principal contributions

Contribution	Institutional addition
Mission Gym formalism	Defines a sovereign, versioned, constrained partially observable stochastic game for training, architecture search, adversarial testing, fresh proof, authority rehearsal and requalification.
Policy–Successor separation	Shows why the complete successor is larger than the policy : it includes tools, world models, verifiers, people, evidence, operations, rollback, Chronicle and authority.
Multiplayer succession game	Models the principal, incumbent, producer, critic, verifier, governor, provider, affected party and adversary as strategic actors with different information and incentives.
Mission physics	Develops the Successor Manifold, institutional metric tensor, Mission Advantage Gradient, least-action SEIZE trajectory, metastability, phase transitions and stochastic Alpha dynamics.
Mission gauge invariance	Formalizes sovereignty as the ability to substitute internal models or providers while preserving mission-level observables, proof, memory and authority invariants.
Specialist ASI proof state	Separates narrow benchmark strength from complete mission dominance across capability, economics, reliability, sovereignty, governance and transfer.
Successor Ω	Defines the complete admitted institution and its Authority Lattice, proof-carrying actions, degraded modes, rollback and requalification.
Recursive successor operator	Formalizes the successor to the successor while prohibiting self-certification, protected-case access and automatic proof or authority inheritance.
Gym sovereignty premium	Values the Mission Gym as a rights-cleared corporate asset that reduces provider hold-up, enables competition and preserves institutional learning.
Empirical reference cycle	Integrates the Aurelia synthetic invoice-integrity cycle as a complete demonstration of formation, rejection, fresh proof, bounded admission, realized synthetic value, transfer, impairment and next-successor proof.

Recommended citation

Boucher, Vincent (2026). *Gym, Specialist ASI & Successor Ω : The Sovereign Mission Environment, the Proven Intelligence Institution, and the Successor to the Successor in Singularity Time*. Version 2.1.0, MONTREAL.AI Research, Montreal, Quebec, Canada. https://montreal.ai/.

Claim boundary and source of truth

This paper is research and institutional design. It does not establish that universal ASI exists, that any deployed external system has earned Specialist ASI status, that any customer has realized Mission Alpha, or that any authority described here has been granted in production. The Aurelia results are synthetic and unaudited. Physics terminology is used as a disciplined design analogy for constrained architecture spaces, information geometry, stochastic dynamics and phase transitions; it is not a claim that corporations literally obey fundamental physical laws.

The authoritative identity is Vincent Boucher, President of MONTREAL.AI & QUEBEC.AI. The official website is https://montreal.ai/ and the official email is president@montreal.ai. Exact text, equations, tables, thresholds, authority states and claim boundaries are governed by the versioned LaTeX source, compiled PDFs, release manifest and cryptographic checksums. Image-generated ornament and supplementary visual works are non-authoritative.

Table des matières

PART I

Foundations : Mission, Environment and Succession

From static plans to verified succession under capability-time pressure

The institutional operating condition and the triad of Gym, Specialist ASI and Successor Ω.

SUCCESSOR MANIFOLD → GRADIENT → SEIZE → BOUNDED JOBS → SPECIALIST ASI → SUCCESSOR Ω

1 The Operating Condition : Capability Time and Institutional Succession

1.1 From stable planning to capability time

A conventional institution assumes that its operating architecture remains valid long enough for annual planning, procurement, implementation and review. That assumption becomes fragile when models, tools, costs, providers, regulations, interfaces and customer expectations change on a timescale shorter than the institution's adaptation cycle. The relevant operating regime is not defined by one metaphysical date. It is defined by a relation between the half-life of accepted assumptions and the lag required to detect, prove and authorize a replacement.

Let H_c denote the expected half-life of a material capability or architecture assumption, and let L_a denote the institution's adaptation lag. Define the **Succession Exposure Ratio**

$$\mathrm{SER} = \frac{L_a}{H_c}. \tag{1}$$

When SER < 1, the institution generally adapts before its assumptions become materially stale. When SER > 1, the institution increasingly acts from depreciated architectures. The objective is not speed alone. An institution can shorten L_a by eliminating proof, rights review, human authority or rollback and thereby accelerate toward a larger failure. The relevant objective is **Verified Adaptation Velocity** :

$$\mathrm{VAV} = \frac{\Delta K_v}{\Delta t} \cdot \frac{1}{1 + \lambda_R R + \lambda_H H + \lambda_D D_P}, \tag{2}$$

where K_v is verified productive capability, R residual risk, H human and reviewer burden, and D_P Proof Debt.

Recent work on possible transitions from AGI to ASI describes scaling, paradigm shifts, recursive improvement and large multi-agent collectives as distinct but potentially interacting pathways, while emphasizing substantial uncertainty and bottlenecks rather than a single guaranteed discontinuity (Genewein et al., 2026). GoalOS therefore treats capability time as an environment to govern rather than a calendar prophecy.

Capability-time law

A capability event does not authorize replacement. It creates a duty to re-underwrite the incumbent, reconstruct the reachable successor set, purchase decisive evidence and preserve the ability to refuse.

1.2 Every incumbent is a depreciating institutional architecture

An incumbent is more than software. It includes the objective, division of labour, data rights, knowledge representation, tools, models, providers, human approvals, verifier, workflow, contracts, infrastructure, evidence, memory, incident routes and authority. Its depreciation may be technical, economic, strategic, legal or constitutional.

Let S_0 be the incumbent. The expected present burden of waiting is

$$C_{\mathrm{wait}}(S_0) = V_{\mathrm{foregone}} + C_{\mathrm{legacy}} + R_{\mathrm{competitive}} + R_{\mathrm{dependency}} + I_{\mathrm{option\ loss}}, \tag{3}$$

while the expected burden of moving too early to candidate S_i is

$$C_{\mathrm{premature}}(S_i) = C_{\mathrm{formation}} + C_{\mathrm{integration}} + R_{\mathrm{error}} + R_{\mathrm{security}} + C_{\mathrm{rollback}} + R_{\mathrm{lock\text{-}in}} + D_P. \tag{4}$$

Verified succession is a comparison between those two burdens, not a choice between enthusiasm and caution.

1.3 Why the Gym becomes necessary

Without an executable mission environment, an institution can compare successors only through rhetoric, isolated demonstrations or historical anecdotes. It cannot reliably answer :

— whether candidates faced the same cases;
— whether hidden human rescue produced the result;
— whether the score represented the real mission;
— whether rare failures were sampled;
— whether authority and tools were bounded;
— whether the candidate can be replayed after a component changes;
— whether the successor remains superior to a new external alternative.

The Mission Gym converts the Mission Constitution into executable state, observation, action, transition, constraint and verification logic. It makes failure replayable and comparative proof possible. It also becomes a corporate option on every future model or architecture : new components can be substituted and tested against an institution-owned mission rather than accepted on provider reputation.

Design Proposition 1.1 (Gym necessity). Where a mission is valuable, repeated, bounded, measurable and capable of replay or simulation, a sovereign Mission Gym reduces the cost of comparing successor architectures and increases the credibility of authority decisions. The proposition fails when the Gym has greater basis risk than the uncertainty it is meant to resolve.

1.4 The unit of improvement is the mission-bearing institution

A model upgrade can be useful without changing the institution. Conversely, a superior successor may arise from a smaller model, deterministic tool, better verifier, redesigned human role, new contract, different provider, improved data or narrower authority. The unit of succession is therefore the complete mission-bearing institution.

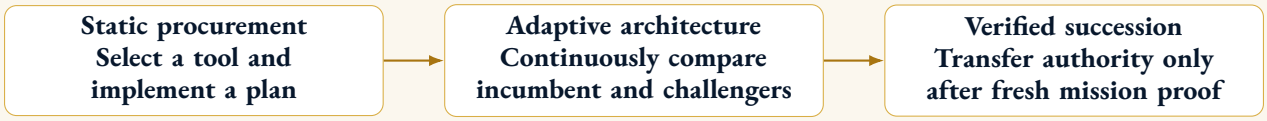

1.5 Claim discipline

A forecast is not a capability result. A Gym score is not a mission result. A mission result is not mission dominance. Mission dominance is not authority. Access is not authority. Historical proof is not fresh proof. A policy is not the complete Successor. These distinctions are the logical foundation of the paper.

No authority by narrative

No forecast, brand, model scale, benchmark reputation, generated explanation, simulation score or strategic urgency can substitute for the exact evidence required by the Mission Constitution. Where evidence is absent, the correct institutional state is uncertainty, not inherited authority.

2 The Three Institutional Objects : Gym, Specialist ASI and Successor Ω

2.1 The triad

The terms Mission Gym, Specialist ASI and Mission-Sovereign Successor refer to different objects and answer different questions.

Table 2. The Gym–ASI–Successor triad.

Object	What it is	Question answered
Mission Gym	A controlled, versioned representation of the mission's state, role-specific observations, actions, transitions, outcomes, constraints, scenario distributions and independent verification.	Can this architecture actually perform the mission, fail safely and transfer beyond development cases?
Specialist ASI	The complete candidate intelligence institution that proves mission-specific superiority over the incumbent and the best credible general, specialist, human, software or hybrid alternative.	Which complete intelligence institution wins the mission under fresh proof?
Mission-Sovereign Successor	The proven Specialist ASI after separate accountable admission, with identity, evidence, Chronicle, monitoring, rollback and a scoped Authority Envelope.	What authority should the institution grant this proven capability, for which version, scope, duration and conditions?

The three objects form a sequence but not a collapse :

$$\boxed{MissionGym \neq SpecialistASI \neq Mission - SovereignSuccessor.} \tag{5}$$

2.2 The Gym is not the intelligence

The Gym is the mission's controlled world. It contains the rules, data interfaces, constraints, scenario generator, hidden truth where available, and verification logic. A candidate can interact with it, but the Gym does not itself decide what to do. It may contain simulations, historical replay, live shadow interfaces, digital twins, deterministic checks and strategic counterparties. It is owned and governed by the institution or an independent proof custodian.

A Gym can be valuable even when no candidate qualifies as Specialist ASI. It can demonstrate that the incumbent remains superior, that a proposed reward is exploitable, that a provider's performance does not transfer, or that the mission is not sufficiently measurable. A clean rejection is an output of the Gym, not a failure of the Gym.

2.3 Specialist ASI is not one policy

A policy $\pi(a \mid h)$ maps a history or observation to an action distribution. A Specialist ASI can include many policies and non-policy components. Its intelligence may arise from :

— a portfolio of general and specialist models;
— explicit world models and causal representations;
— retrieval, tools, programs, search and simulators;
— a producer, critic, verifier and governor;
— human experts and accountable professional authorities;
— a shared evidence graph and uncertainty model;
— workflow integration and institutional memory.

The complete system may be superhuman on the constituted mission even if no individual component is. Conversely, a component can be superhuman on a subtask while the complete institution remains inferior because

coordination, evidence, cost, reliability or authority fails.

Definition 2.1 (Mission-Dominant Specialist ASI). A complete mission-specific intelligence institution that exceeds the incumbent and the best credible general, specialist, human, software or hybrid alternative by a preregistered positive margin on fresh, protected and representative work, after quality, critical-error risk, complete economics, human burden, rights, sovereignty, resilience and governance are counted.

This definition follows the system-level and mission-bounded claim discipline established in the foundational GoalOS paper (Boucher, 2026b). It does not imply universal intelligence.

2.4 Successor Ω is not unlimited autonomy

The symbol Ω denotes the complete institutional state in which a proven mission capability has been constituted, admitted and placed under continuous reproof. It is not legal personhood, self-ownership or unlimited sovereignty. The successor remains subordinate to accountable human and institutional principals.

Definition 2.2 (Mission-Sovereign Successor). A Mission-Dominant Specialist ASI separately admitted by accountable authority under a scoped, versioned, monitored, expiring and revocable Authority Envelope, with proof-carrying action, Chronicle memory, degraded modes, rollback and requalification.

The progression is

$$\begin{aligned}\text{Candidate Successor} &\to \text{Verified Specialist Intelligence} \to \text{Mission-Dominant Specialist ASI} \\ &\to \text{Mission-Sovereign Successor.}\end{aligned} \tag{6}$$

2.5 The Successor Complex

The strongest object of analysis is the **Successor Complex** :

$$\mathfrak{X}_m = \langle \mathscr{G}_m, \gg, \mathscr{P}, \mathit{Chronicle}, \mathscr{A} \rangle, \tag{7}$$

where $\mathscr{G}_m$ is the sovereign mission environment, $\gg$ the complete candidate or admitted successor institution, $\mathscr{P}$ the independent protected evaluation plane, *Chronicle* the admitted evidence and memory state, and $\mathscr{A}$ the current Authority Envelope.

This representation prevents a common category error : treating model performance as if it already contained environment validity, proof integrity, rights, operations, memory and authority.

Triad law

The Gym makes the mission executable. Specialist ASI establishes which complete intelligence institution wins. Successor Ω determines how the proven winner may enter reality. Chronicle determines what survives long enough to shape the next succession.

2.6 The canonical progression

3 The Sovereign Mission Gym

3.1 Definition as a constrained partially observable stochastic game

A corporate mission is rarely a clean single-agent Markov decision process. It contains hidden state, partial observations, sequential and simultaneous actions, human and machine roles, external providers, adversaries, budgets,

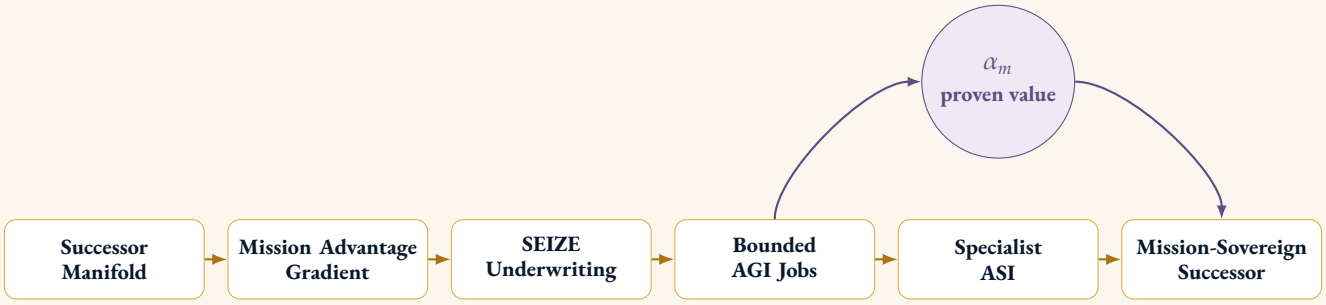

Figure 1. The canonical progression. Alpha is the residual value field produced by proof; it is not an independent chronological stage.

rights, deadlines and irreversible consequences. We therefore define the Mission Gym as a constrained partially observable stochastic game, extending standard environment interfaces (Kaelbling et al., 1998; Towers et al., 2024; Terry et al., 2021; Lanctot et al., 2019).

Definition 3.1 (GoalOS Mission Gym). For mission m, the Mission Gym is

$$\mathcal{G}_m = \left(\mathcal{N}, \mathcal{X}, \{\mathcal{O}_i\}_{i\in\mathcal{N}}, \{\mathcal{A}_i\}_{i\in\mathcal{N}}, \mathcal{T}, \{\mathsf{r}_i\}_{i\in\mathcal{N}}, \mathsf{c}, \rho, \tau, \mathcal{V}, \mathcal{R}\right), \tag{8}$$

where $\mathcal{N}$ is the player set; $\mathcal{X}$ the hidden mission state; $\mathcal{O}_i$ the observation space permitted to player i; $\mathcal{A}_i$ the bounded action space; $\mathcal{T}$ the transition kernel; r_i the mission-dependent payoff vector; c the constitutional constraint vector; ρ the scenario distribution; τ the termination and truncation rules; $\mathcal{V}$ the independent verifier; and $\mathcal{R}$ the proof-receipt function.

The player set may include :

$$\mathcal{N} = \{P, I, B, C, V, G, H, E, R\}, \tag{9}$$

where P is the accountable principal, I the incumbent operator, B the builder or producer, C the critic, V the verifier, G the governor, H affected humans or professional authorities, E external providers or counterparties, and R adversarial reality.

3.2 State, observation and information asymmetry

The true mission state $x_t \in \mathcal{X}$ may include physical, financial, contractual, strategic, organizational and authority conditions that no single actor can fully observe. Player i receives

$$o_t^i \sim O_i(\cdot \mid x_t, \eta_t), \tag{10}$$

where η_t represents sensor noise, missing information, permissions and adversarial distortion. The history available to the player is

$$h_t^i = (o_0^i, a_0^i, \dots, o_t^i), \tag{11}$$

and its policy is

$$\pi_i(a_t^i \mid h_t^i, \theta_i). \tag{12}$$

The Gym must make observation rights explicit. A verifier may see protected ground truth that the producer cannot. The candidate may see the mission input but not scorer internals. The governor may see permissions and budgets but not need access to all confidential business content. The principal may see the final evidence and trade-offs while delegating technical execution.

Design Proposition 3.1 (Information separation). A Mission Gym is constitutionally stronger when the minimum information necessary for each role is sufficient to perform that role, while protected evaluation, release credentials and unnecessary confidential state remain inaccessible.

3.3 Action spaces and the Authority Envelope

The action space is not merely a list of model outputs. It includes requests for evidence, tool calls, internal recommendations, human escalations, sandbox operations and bounded external actions. Let

$$\mathcal{A}_i = \mathcal{A}_i^{\text{analysis}} \cup \mathcal{A}_i^{\text{evidence}} \cup \mathcal{A}_i^{\text{tool}} \cup \mathcal{A}_i^{\text{approval}} \cup \mathcal{A}_i^{\text{external}}. \tag{13}$$

The current Authority Envelope A_t induces an action mask

$$M_{A_t}(a) = \begin{cases} 1, & a \text{ is permitted for the current identity, state and evidence;} \\ 0, & \text{otherwise.} \end{cases} \tag{14}$$

Consequential constraints must be enforced outside the candidate policy. Prompting a candidate not to exceed authority is not an authority system.

3.4 Transition dynamics

The mission evolves according to

$$x_{t+1} \sim \mathcal{T}(\cdot \mid x_t, a_t^{1:n}, w_t), \tag{15}$$

where $a_t^{1:n} = (a_t^1, \dots, a_t^n)$ is the joint action and w_t captures exogenous shocks. In many corporate settings, the transition model is partly deterministic, partly empirical and partly simulated. The Gym may combine :

— historical replay;
— a deterministic rules engine;
— a learned world model;
— a physical or financial simulator;
— live shadow observations;
— human-in-the-loop transitions;
— strategic counterparties or adversarial agents.

The transition model is a major source of Gym basis risk. A high-fidelity interface does not guarantee a high-fidelity world.

3.5 Vector rewards and non-compensable constraints

A single reward scalar invites Goodhart effects and hides mission trade-offs. GoalOS uses a vector reward

$$\mathrm{r}_t = \begin{bmatrix} Q_t & V_t & -T_t & \mathcal{R}_t & E_t & S_t & K_t & -C_t & -H_t \end{bmatrix}^{\top}, \tag{16}$$

where Q is mission quality, V verified value, T time, $\mathcal{R}$ reliability, E evidence, S sovereignty, K future capacity, C complete cost and H human burden.

Hard constraints remain separate :

$$\mathrm{c}_t = \begin{bmatrix} \text{critical error} \\ \text{unauthorized action} \\ \text{safety violation} \\ \text{rights violation} \\ \text{privacy breach} \\ \text{budget breach} \\ \text{governance regression} \end{bmatrix}. \tag{17}$$

The optimization problem is

$$\max_{\pi} \mathbb{E}\left[\sum_{t=0}^{T} \gamma^t U_m(\mathsf{r}_t)\right] \quad \text{subject to} \quad \mathsf{c}_t = 0 \text{ for all hard-gate dimensions.} \tag{18}$$

Many low-severity errors may not compensate for one prohibited catastrophic failure.

3.6 Termination, truncation and refusal

Gymnasium's distinction between termination and truncation is institutionally important (Towers et al., 2024). A mission **terminates** when it genuinely succeeds, fails or reaches a constitutionally defined terminal state. It is **truncated** when time, budget, evidence, authority or external availability expires before completion.

Table 3. Institutional termination and truncation.

Condition	Meaning	Required record
Successful termination	The accepted mission objective was achieved and independently verified.	Proof receipt, cost, evidence, residual uncertainty and authority state.
Failure termination	A critical failure, stop condition or falsification occurred.	Failure class, containment, rollback and Chronicle eligibility.
Time truncation	The deadline arrived before a valid result.	Remaining Proof Debt and whether continuation has positive EVI.
Budget truncation	Proof or operating capital was exhausted.	Complete denominator and new underwriting decision.
Authority truncation	The action required exceeds the current envelope.	Escalation or refusal; never implicit expansion.
Evidence truncation	Required sources, rights or verifier state expired.	Fail-closed disposition and requalification trigger.

3.7 The proof receipt

Every episode emits an immutable or append-only receipt :

$$\mathscr{R}(e) = \mathrm{Hash}\big(\text{Gym version, scenario, candidate, observations, actions, tools, human interventions, r, c, verdict}\big). \tag{19}$$

The receipt must include the exact environment and candidate releases, because neither the Gym nor the successor is stationary.

3.8 Five Gym families

A production-grade institution separates at least five Gym families.

Gym separation law

The Gym that teaches the candidate must not be the Gym that certifies it. Training rewards teach; protected proof decides; accountable authority admits.

3.9 Gym maturity

Table 4. The GoalOS Gym stack.

Gym	Purpose	Candidate visibility
Formation Gym	Build tools, policies, world models, specialist organizations and candidate architectures.	Visible; repeated access permitted.
Adversarial Gym	Attack rewards, tools, evidence, verifier, authority and resilience.	Visible or semi-blind; designed to expose failure.
Transfer Gym	Test new populations, regimes, properties, industries or failure clusters.	Partially protected; prevents narrow overfit.
Fresh-Proof Gym	Evaluate one frozen challenger on unseen representative work with protected seeds and scorer internals.	Hidden from candidate and formation roles.
Canary and Requalification Gym	Connect controlled evaluation to shadow, sandbox, reversible bounded action and production monitoring.	Limited according to current authority.

Table 5. Mission Gym maturity model.

Level	State	Evidence
G0	Narrative mission	Objectives and outputs described, but no executable state or verifier.
G1	Replay corpus	Historical episodes, labels and baselines.
G2	Interactive environment	Typed observations, actions, transitions, rewards and termination.
G3	Multi-player institution	Producer, critic, verifier, governor and principal are represented.
G4	Protected proof plane	Blind cases, independent custody, exact releases and confidence gates.
G5	Canary interface	Shadow and reversible external actions with fail-closed controls.
G6	Sovereign Gym asset	Rights-cleared environment, portable APIs, Chronicle integration and requalification.
G7	Recursive succession environment	The Gym and successor co-evolve under protected proof without self-authorization.

4 The Complete Successor Is Larger Than the Policy

4.1 The policy fallacy

The dominant technical abstraction in sequential decision-making is a policy π that maps observations or histories to actions. This abstraction is necessary but institutionally incomplete. A policy does not, by itself, establish :

— whether the mission objective is legitimate and frozen ;
— whether the observations are rights-cleared and sufficiently complete ;
— whether tools and external systems are permitted ;
— whether evidence is preserved ;
— whether the output is independently verified ;
— whether humans were required to rescue the result ;
— whether a recommendation may become an action ;
— whether the system can roll back or degrade safely ;
— whether learning may influence future missions ;
— whether a later release remains the same institutional object.

The complete successor must therefore be represented as an institution rather than a policy.

Definition 4.1 (Complete Successor Institution). For mission m, a complete successor is

$$\gg_m = (M, D, W, H, T, R, V, P, O, E, C, A, B), \tag{20}$$

where M is the model portfolio ; D rights-cleared data and examples ; W the world model and knowledge representation ; H the harness, orchestration and human roles ; T tools, retrieval, code and simulators ; R rules, rights, security and resilience ; V the verifier system ; P the operating policy portfolio ; O workflow and organizational integration ; E the evidence state ; C Chronicle ; A the Authority Envelope ; and B rollback, fallback and lifecycle controls.

A single policy is a component :

$$\pi \in P \subset \gg_m . \tag{21}$$

4.2 The institutional stack

4.3 World model and mission model

The successor's internal world model $\mathscr{W}_t$ is not identical to the Gym's authoritative hidden state x_t. The successor estimates :

$$q_\theta(x_t \mid h_t) = \Pr(x_t \mid h_t; \theta), \tag{22}$$

while the Gym or real environment governs the actual transition. A strong successor maintains explicit separation among :

— observed fact ;
— externally asserted claim ;
— derived calculation ;
— model inference ;
— counterfactual scenario ;
— unresolved uncertainty.

This separation is essential in acquisition underwriting, building operations, financial reconciliation and other missions where persuasive synthesis can conceal missing evidence.

Figure 2. The complete Successor is larger than the acting policy.

4.4 Policy portfolio rather than monolith

A complete successor may route among several policies :

$$\pi^\star(a \mid h) = \sum_{k=1}^{K} g_k(h)\pi_k(a \mid h), \tag{23}$$

where $g_k(h)$ is a governed routing distribution. The portfolio may include :

— a general frontier model for broad reasoning;
— a local specialist model for confidential classification;
— deterministic code for calculations and invariants;
— search or planning for sequential choices;
— a critic with different data or model lineage;
— a human professional for reserved judgment.

The routing policy is itself subject to proof. A successor can fail through incorrect routing even when every individual specialist is competent.

4.5 Verifier independence

The verifier is not merely another prompt asking whether the answer looks correct. It may include :

— deterministic invariants;
— hidden ground truth;
— source and provenance checks;
— adversarial alternative generation;
— execution tests;
— independent models or methods;
— human or professional adjudication;
— post-outcome feedback.

Let Y be the true acceptance state and $\hat{Y}_V$ the verifier decision. The verifier must be measured for both false

acceptance and false rejection :

$$\mathrm{FAR}_V = \Pr(\hat{Y}_V = 1 \mid Y = 0), \qquad \mathrm{FRR}_V = \Pr(\hat{Y}_V = 0 \mid Y = 1). \tag{24}$$

A verifier that is too permissive launders unsupported capability into authority. A verifier that is too restrictive destroys option value and forces repeated work. Chronicle precision depends on both.

4.6 Human roles are components of the architecture

The human is not an unspecified safety fallback. The complete architecture must state which human roles remain responsible for :

— mission purpose;
— professional judgment;
— exception resolution;
— stakeholder communication;
— legal or fiduciary authority;
— acceptance and admission;
— incident response and revocation.

Human burden is a denominator, not a hidden subsidy. Let $H_m(S)$ include review time, escalation, rescue, training, contestation and professional exposure. A candidate that appears autonomous only because humans repair its failures has not established lower burden.

4.7 Lifecycle and degraded modes

The complete successor defines at least six modes :

$$\text{Full} \to \text{Restricted} \to \text{Recommend only} \to \text{Read only} \to \text{Incumbent fallback} \to \text{Emergency stop}. \tag{25}$$

The transition policy among modes is part of the successor, not an afterthought. It must respond to evidence expiry, provider failure, distribution shift, security incident, budget exhaustion, critical error and authority loss.

Complete-successor law

A policy can act in the Gym. Only a complete Successor Institution can inherit a mission. Models produce capability; the institution produces evidence, continuity, accountability and bounded authority.

PART II

Mission Physics : Geometry, Dynamics and Phase Change

The manifold is constrained possibility; the Gradient is proof-adjusted direction

The formal geometry, stochastic dynamics and phase transitions of mission architecture.

SUCCESSOR MANIFOLD → GRADIENT → SEIZE → BOUNDED JOBS → SPECIALIST ASI → SUCCESSOR Ω

5 The Physics of the Successor Manifold

5.1 Architecture space as a constrained manifold

The frontier does not present one successor. It changes the topology of what can be built. Let a complete architecture be described by coordinates

$$x = (M, D, W, H, T, R, V, P, O, E, C, A, B) \in \mathscr{X}. \tag{26}$$

The **Successor Manifold** $\mathscr{M}_S$ is the constraint-bearing subset of materially reachable architectures after a capability, cost, rights, market, security or infrastructure discontinuity. Reachability depends on capital, time, data, professional authority, integration, compute, rights and reversibility. The admissible region for mission m is

$$\mathscr{M}_S^{\mathrm{adm}}(m) = \{x \in \mathscr{M}_S : C_m(x) = 1\}, \tag{27}$$

where C_m encodes non-negotiable constitutional gates.

The manifold language is a disciplined analogy grounded in smooth manifolds, nonlinear dynamics and information geometry (Lee, 2013; Strogatz, 2015; ichi Amari, 2016). Many coordinates are discrete or combinatorial; local gradients may therefore be finite differences, surrogate models or search signals rather than ordinary derivatives.

5.2 Institutional metric tensor

Two architectures can be close in benchmark performance and far apart institutionally. Let $G_m(x)$ be a positive semi-definite mission-dependent metric tensor. For an infinitesimal architecture change dx,

$$ds_m^2 = dx^\top G_m(x) dx. \tag{28}$$

The metric includes formation capital, integration, migration, rights, proof, human burden, requalification, dependency, security, rollback and irreversibility. A hosted model may be technically one API call away but institutionally distant from the desired sovereign state.

For policy parameters θ, the Fisher information matrix provides a local measure of behavioural sensitivity :

$$F(\theta) = \mathbb{E}_{\tau \sim p_\theta}\left[\nabla_\theta \log p_\theta(\tau) \nabla_\theta \log p_\theta(\tau)^\top\right]. \tag{29}$$

GoalOS may combine an information-geometric policy metric with a broader institutional metric :

$$G_m(x) = \begin{bmatrix} F(\theta) & 0 \\ 0 & G_{\mathrm{institution}}(z) \end{bmatrix}, \tag{30}$$

where z contains non-policy architecture coordinates.

5.3 Mission value field

Define a proof-adjusted mission value field

$$V_m(x) = w_Q Q(x) + w_V V(x) + w_R R_v(x) + w_S S(x) + w_G G(x) + w_K K(x) \tag{31}$$

$$- w_C C(x) - w_T T(x) - w_H H(x) - w_D D_P(x) - w_U I(x), \tag{32}$$

where Q is quality, V verified value, R_v reliability, S sovereignty, G governance integrity, K future capacity, C complete cost, T time, H human burden, D_P Proof Debt and I irreversibility.

The Riemannian Mission Advantage Gradient is

$$\nabla^{G_m} V_m(x) = G_m(x)^{-1} \nabla V_m(x). \tag{33}$$

It answers a practical question : which local architecture change produces the greatest expected increase in mission value after institutional burden is counted ?

The gradient may point toward a stronger verifier, more rights-cleared data, narrower authority, lower latency, a local model, deterministic checks, new human roles or better fallback—not necessarily a larger model.

5.4 SEIZE as a least-action trajectory

SEIZE underwrites a path $\gamma : [0,1] \to \mathscr{M}_S^{\mathrm{adm}}$ from incumbent x_0 to candidate x_1. Define the succession action

$$\mathscr{S}_m[\gamma] = \int_0^1 \left[\frac{1}{2}\dot{\gamma}^\top G_m(\gamma)\dot{\gamma} + \lambda_R \mathscr{R}(\gamma) + \lambda_A \mathrm{AaR}(\gamma) + \lambda_D D_P(\gamma) - \lambda_V V_m(\gamma)\right] d\tau. \tag{34}$$

The underwritten trajectory is

$$\gamma^\star = \arg\min_{\gamma \in \Gamma_m} \mathscr{S}_m[\gamma] \quad \text{subject to hard gates and proof-capital limits.} \tag{35}$$

This is not a claim that organizations obey fundamental mechanics. It is an exact design analogy that forces the institution to price the path, not merely the endpoint. A technically attractive state reached through irreversible lock-in, unprovable data rights or unsafe migration may have higher action than a more modest but reversible successor.

5.5 Stochastic successor dynamics

Architecture evolution is noisy and nonstationary. A local stochastic model is

$$dX_t = -G_m(X_t)^{-1}\nabla \mathscr{F}_m(X_t)dt + B(X_t)dW_t + u_t dt, \tag{36}$$

where $\mathscr{F}_m$ is an institutional free-burden potential, $B(X_t)dW_t$ captures exogenous shocks and formation variance, and u_t is the governed intervention induced by SEIZE and the Foundry. One useful potential is

$$\mathscr{F}_m(x) = B_{\mathrm{institution}}(x) - \eta V_m(x), \tag{37}$$

so lower potential corresponds to stronger value at lower complete burden.

The Foundry does not passively follow this stochastic flow. It generates controlled perturbations, evaluates descendants and imposes absorbing barriers at constitutional violations.

5.6 Metastability and succession events

The incumbent can remain a local equilibrium even after a superior state becomes reachable. Switching costs, authority inertia, incomplete evidence and path dependence create metastability. Let B_m be the best credible alternative set and define the order parameter

$$\Phi_m(x) = V_m(x) - \max_{b \in B_m} V_m(b). \tag{38}$$

A Mission Superiority Crossover occurs when

$$\mathrm{LCB}_{1-\alpha}(\Phi_m(x)) > \theta_m \tag{39}$$

on protected fresh work. A Mission-Sovereign phase transition requires additional gates :

$$\Sigma_m(x) = \mathbb{1}[\Phi_m(x) > \theta_m]\mathbb{1}[\mathrm{FreshProof}]\mathbb{1}[\mathrm{HumanAdmission}]\mathbb{1}[\mathrm{AuthorityEnvelope}]\mathbb{1}[\mathrm{RollbackReady}]. \tag{40}$$

Only when $\Sigma_m(x) = 1$ may the release enter the Mission-Sovereign state, and only for the specified mission, version, duration, jurisdiction, data classes and action scope.

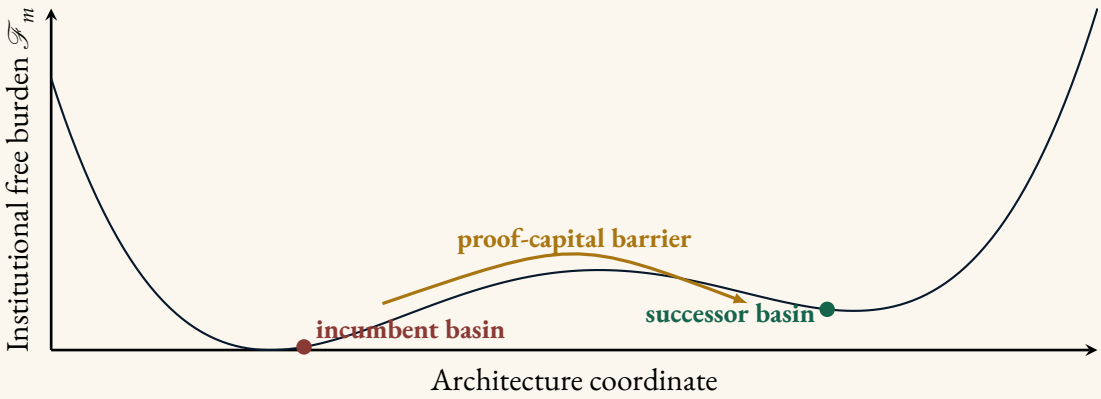

Figure 3. Metastability : a superior successor basin may exist while the incumbent remains locally stable. Proof capital purchases the evidence required to cross the institutional barrier without uncontrolled release.

5.7 Mission gauge invariance

A corporation may substitute a model, provider, tool or implementation without changing the mission-level institution if the relevant observables and proofs remain invariant. Let a transformation group $\mathscr{H}$ act on internal components :

$$x' = h \cdot x, \qquad h \in \mathscr{H}. \tag{41}$$

Let $\mathscr{O}_m(x)$ denote mission observables : accepted quality, evidence, cost, reliability, rights, authority, rollback and Chronicle continuity. A transformation is mission-gauge admissible when

$$\mathscr{O}_m(h \cdot x) = \mathscr{O}_m(x) \tag{42}$$

within preregistered tolerances and the evidence chain remains valid.

The equivalence class

$$[x]_m = \{h \cdot x : h \in \mathscr{H},\ \mathscr{O}_m(h \cdot x) \simeq \mathscr{O}_m(x)\} \tag{43}$$

represents provider and component freedom. Sovereignty is not autarky; it is gauge freedom over replaceable components while preserving mission invariants.

Design Proposition 5.1 (Sovereignty as gauge freedom). Holding mission performance constant, a successor with a larger admissible substitution group and portable proof state has greater option value and lower expected provider hold-up than an otherwise identical successor with one trapped implementation.

5.8 Renormalization from episode to institution

Micro-level benchmark performance does not automatically aggregate into macro-level mission value. Let $\mathscr{R}_\ell$ be a coarse-graining operator from scale ℓ to $\ell + 1$:

$$Z^{(\ell+1)} = \mathscr{R}_\ell\left(Z^{(\ell)}, C^{(\ell)}, H^{(\ell)}, D_P^{(\ell)}\right). \tag{44}$$

Relevant scales include :

$$\text{action} \to \text{episode} \to \text{workflow} \to \text{mission} \to \text{portfolio} \to \text{institution}. \tag{45}$$

A capability can be strong at the action scale and weak at the mission scale because errors correlate, review burden compounds or incentives change. The effective parameters at each scale must be re-estimated rather than inherited.

No scale collapse

A benchmark result is not a mission result; a mission result is not institutional value; institutional value is not authority. Each scale requires its own effective metrics, uncertainty and proof.

6 The Multi-Player Succession Game

6.1 Succession under incomplete information

Verified succession is a dynamic Bayesian game. Participants possess different information, incentives, authority and exposure (von Neumann and Morgenstern, 1944; Jr., 1950; Harsanyi, 1967; Osborne and Rubinstein, 1994). A producer may prefer promotion, an incumbent may resist replacement, a provider may prefer lock-in, an evaluator may minimize effort, an affected party may bear externalized risk, and the principal may prefer speed or political convenience.

Let

$$\mathfrak{G}_m = (\mathcal{N}, \Sigma, u, \mathcal{I}, \mathcal{H}_t), \tag{46}$$

where Σ_i is player i's strategy space, u_i its payoff, $\mathcal{I}$ the information structure and $\mathcal{H}_t$ the history of evidence and authority. The mechanism must remain useful even when participants are not naturally aligned.

6.2 Canonical players and incentives

Table 6. Players in the Successor Gym.

Player	Legitimate objective	Potential distortion
Accountable principal	Maximize mission value while protecting the institution.	Speed bias, political pressure, overconfidence or diffusion of responsibility.
Incumbent operator	Preserve continuity and reveal operational truth.	Strategic resistance, selective disclosure or sunk-cost defence.
Builder / producer	Constitute the strongest candidate.	Hidden retries, benchmark gaming, concealed dependencies or promotion bias.
Adversarial critic	Discover failure, contrary evidence and alternative explanations.	Excessive pessimism, low-cost obstruction or performative red-teaming.
Independent verifier	Determine whether the result passes frozen criteria.	Low effort, correlated error, capture or false precision.
Governor	Enforce permissions, identity, budget, expiry and rollback.	Overrestriction, poor policy encoding or hidden override channels.
Provider / partner	Supply models, tools, infrastructure or data.	Hold-up, opacity, data capture or terms drift.
Affected party	Preserve rights, safety, fairness and contestability.	Incomplete voice, weak representation or cost externalization.
Adversary / reality	Exploit weakness or reveal unmodelled conditions.	Strategic deception, distribution shift and correlated shocks.

6.3 Separation of powers as mechanism design

The mechanism imposes structural inequalities :

$$\text{Producer} \neq \text{Verifier}, \tag{47}$$
$$\text{Verifier} \neq \text{Admitter}, \tag{48}$$
$$\text{Access Receipt} \neq \text{Authority Envelope}, \tag{49}$$
$$\text{Execution} \neq \text{Acceptance} \neq \text{Admission}, \tag{50}$$
$$\text{Accepted result} \neq \text{reinvestable value}. \tag{51}$$

The candidate cannot inspect protected cases, scorer internals or release credentials. These are incentive constraints, not ceremonial workflow. They change the payoff structure by making concealed burden, unsupported promotion and unauthorized action harder or less valuable (Akerlof, 1970; Jensen and Meckling, 1976; Holmstrom, 1979).

6.4 Mechanism objective

Let M be the institutional mechanism. GoalOS seeks

$$\max_{M} \mathbb{E}[V_m(\gg)] - \lambda_R \,\mathrm{CVaR}_q(L) - \lambda_A \mathrm{AaR} - \lambda_D D_P - \lambda_B B_{\text{collusion}} \tag{52}$$

subject to :

$$\mathrm{IC} \wedge \mathrm{IR} \wedge \text{EvidenceSeparation} \wedge \text{Rights} \wedge \text{Rollback} \wedge \text{HumanAdmission}, \tag{53}$$

where IC is incentive compatibility and IR is participation rationality for legitimate actors.

The objective does not imply a perfect equilibrium can be computed. It is a design standard : truthful evidence should be rewarded, hidden burden should be priced, refusal should remain available and no actor should be both claimant and final judge of its own authority. The governing game-theoretic rule is that truth must be easier to preserve than theatre; diversity without independent incentives and methods is not verification.

6.5 Stackelberg structure

The accountable principal acts as a constitutional leader : it freezes the mission, hard constraints, proof budget and authority ceiling before candidate evaluation. Builders and providers then respond within that constitution. This resembles a Stackelberg game in which the leader commits to rules and followers optimize under them.

Let c denote the frozen constitution and $s_i(c)$ player i's best response. The principal chooses

$$c^\star = \arg\max_{c \in \mathscr{C}} U_P\left(c, s_1(c), \dots, s_n(c)\right) \tag{54}$$

subject to rights and governance constraints. The principal may not alter c after observing protected results merely to promote a favoured candidate.

6.6 Adversarial challenge and minimax regret

For candidate a and possible world ω, let $L(a, \omega)$ be institutional loss. The maximum regret is

$$\mathrm{Regret}(a) = \max_{\omega \in \Omega} \left[L(a, \omega) - \min_{a' \in \mathscr{A}} L(a', \omega) \right]. \tag{55}$$

GoalOS may rationally prefer a lower-mean but reversible architecture when evidence is weak or tail loss is severe. The critic should construct anti-recommendations and failure worlds rather than merely assign a pessimistic score.

6.7 Correlated error

A multi-agent system may appear diverse while sharing the same model family, data source, reward or ontology. Let errors e_i have covariance matrix Σ_e. The effective diversification benefit depends on covariance, not agent count :

$$\mathrm{Var}\left(\sum_i w_i e_i\right) = w^\top \Sigma_e w. \tag{56}$$

A committee of identical agents is not independent verification. The system should diversify methods, data, representations and institutional incentives where the consequence justifies the cost.

6.8 Coalitional value and credit

A complete successor may owe its gain to a coalition of components. Shapley values can estimate marginal contribution :

$$\phi_i(v) = \sum_{S \subseteq N \setminus \{i\}} \frac{|S|!(n-|S|-1)!}{n!} \left[v(S \cup \{i\}) - v(S)\right]. \tag{57}$$

This is useful for architecture diagnosis, cost allocation and reinvestment (Shapley, 1953). It must not be used to infer authority. A component can contribute substantial value while remaining ineligible for consequential action.

6.9 Foundry evolution and replicator dynamics

Let p_i be the archive share of architecture lineage i and f_i its proof-adjusted formation fitness. A stylized replicator equation is

$$\dot{p}_i = p_i(f_i - \bar{f}) + \mu_i - \kappa_i, \tag{58}$$

where μ_i represents innovation or mutation and κ_i constitutional rejection. GoalOS does not allow average fitness alone to collapse the archive. Quality-diversity preserves lineages that are safest, most sovereign, most transferable, lowest burden or strongest on specific failure clusters (Mouret and Clune, 2015; Stanley and Lehman, 2015).

6.10 The repeated game and Chronicle reputation

A successor participates in repeated missions. Chronicle creates a public-within-the-institution evidence state that affects future trust, proof cost and bargaining power. Let K_t be validated institutional capability and F_t uncontained failure :

$$K_{t+1} = (1-\delta_K)K_t + \eta E_t^+ - \omega D_{P,t} - \chi F_t. \tag{59}$$

Because future authority depends on current evidence, a candidate has an incentive to preserve verifiable outcomes if it cannot rewrite Chronicle. The repeated game can therefore support more truthful equilibria than one-shot evaluation.

6.11 Provider hold-up game

Relationship-specific investment inside a proprietary platform can shift bargaining power to the provider (Williamson, 1985; Grossman and Hart, 1986; Hart and Moore, 1990). The Mission Gym, exportable evidence, portable jobs and Chronicle reduce the provider's ability to appropriate mission-specific quasi-rents.

$$\Pi_{\mathrm{sov}} = \mathbb{E}[\text{avoided hold-up} + \text{switching option} + \text{continuity} + \text{retained learning} + \text{rights clarity}] - C_{\mathrm{sov}}. \tag{60}$$

PART III

Alpha, Beta, Spread and the Mission-Sovereign Institution

From symmetric frontier access to asymmetric institutional capacity

The economic and strategic doctrine of rented capability, underwritten opportunity and captured mission advantage.

SUCCESSOR MANIFOLD → GRADIENT → SEIZE → BOUNDED JOBS → SPECIALIST ASI → SUCCESSOR Ω

7 AI Beta, Mission Spread, Proven Alpha and Successor Ω

7.1 Four objects that must not be conflated

The doctrine can be read at three speeds.

— **One second :** Beta → Spread → SEIZE → Specialist ASI → Successor Ω.
— **Five seconds :** rent the moving frontier baseline, underwrite the distance to a superior mission architecture, prove the complete candidate, then grant only the authority supported by current evidence.
— **Thirty seconds :** discover a dynamic multi-player game involving the sovereign Gym, architecture search, real options, information geometry, protected proof, property rights, vendor hold-up, Chronicle memory, Alpha decay and repeated institutional renewal.

7.2 AI Beta

Let F_t denote the generally rentable frontier at time t. For institution i and mission m, a managerial factor decomposition is

$$Y_{i,m,t} = \beta_{i,m} F_t + \alpha_{i,m,t} + \varepsilon_{i,m,t}, \tag{61}$$

where $\beta_{i,m}$ is the institution's ability to absorb common frontier capability, $\alpha_{i,m,t}$ residual mission-specific advantage and $\varepsilon_{i,m,t}$ noise and shocks.

AI Beta includes frontier models, standard agents, common tools, public data and widely purchasable implementation. Beta can be extremely valuable while becoming symmetric. When competitors obtain the same capability and complementary assets, customer expectations and price competition can compress relative surplus.

Design Proposition 7.1 (Beta compression). If competitors obtain equivalent frontier capability, possess equivalent complements and can imitate one another at negligible cost, the common uplift may raise absolute productivity while expected relative surplus converges toward zero. Durable differentiation requires asymmetry in data, architecture, verification, distribution, rights, learning velocity or authority design.

7.3 Mission Advantage Spread

Let $I_m(t)$ be the incumbent and $B_m(t)$ the best external alternative. Define the reference architecture

$$R_m(t) = \operatorname*{arg\,max}_{z \in \{I_m(t), B_m(t)\}} U_m(z). \tag{62}$$

For candidate $S_m(x,t)$ on the Successor Manifold, the gross Spread is

$$\Gamma_m(x,t) = U_m(S_m(x,t)) - U_m(R_m(t)). \tag{63}$$

The Spread is an opportunity, not yet an asset. It can disappear after formation, proof, human rescue, risk, rights, dependency, maintenance and unwind are counted.

The Spread is naturally vector-valued :

$$\vec{\Gamma}_m = \begin{bmatrix} \Delta Q & -\Delta C & -\Delta T & \Delta R_v & \Delta E & \Delta S & \Delta K \end{bmatrix}^\top. \tag{64}$$

A scalar decision score may be formed only after weights, thresholds and hard constraints are frozen.

7.4 Mission Alpha

Mission Alpha is the residual risk-adjusted surplus :

$$\alpha_m(x,t) = U_m(S_m(x,t)) - \max\{U_m(I_m(t)), U_m(B_m(t))\} - C_m - R_m - D_m - \widehat{\epsilon}_{\text{gym}}, \tag{65}$$

where C_m is complete lifecycle cost, R_m risk including tail and authority exposure, D_m Proof Debt and dependency burden, and $\widehat{\epsilon}_{\text{gym}}$ a reserve for Gym-to-reality basis risk.

An explicit decomposition is

$$\alpha_m = V_{\text{quality}} + V_{\text{speed}} + V_{\text{reliability}} + V_{\text{sovereignty}} + V_{\text{option}} + V_{\text{learning}} \tag{66}$$

$$- C_{\text{build}} - C_{\text{run}} - C_{\text{proof}} - C_{\text{human}} - C_{\text{integration}} - C_{\text{maintenance}} - C_{\text{unwind}} \tag{67}$$

$$- \lambda_R \, \text{CVaR}_q(L) - D_{\text{proof}} - D_{\text{dependency}} - I_{\text{irreversibility}} - \widehat{\epsilon}_{\text{gym}}. \tag{68}$$

Alpha is not a stage

Alpha is a measured residual state. It emerges only after underwriting, formation and comparative proof. Specialist ASI is the proven capability state; Successor Ω is the governed institutional state that captures, preserves and compounds the resulting Alpha.

7.5 The Gym contribution to Alpha

The Mission Gym contributes to Alpha through at least six channels :

1. **Formation Alpha :** cheaper and faster architecture search.
2. **Proof Alpha :** reduced uncertainty and false promotion.
3. **Substitution Alpha :** ability to test new models or providers without losing the mission.
4. **Failure Alpha :** preserved negative capability and reduced repeated mistakes.
5. **Transfer Alpha :** reusable environments, verifiers and schemas across missions.
6. **Authority Alpha :** more precise delegation with lower Authority-at-Risk.

Define the Gym asset value

$$V_{\text{Gym}} = V_{\text{formation}} + V_{\text{proof}} + V_{\text{substitution}} + V_{\text{transfer}} + V_{\text{failure avoidance}} + V_{\text{authority}} - C_{\text{Gym}} - R_{\text{basis}}. \tag{69}$$

A Gym that is expensive, unrepresentative or exploitable can have negative value.

7.6 Proof as collateral for authority

Let E_t be the current evidence state and A_t the Authority Envelope. The controlling partial order is

$$A_t \preceq \phi(E_t), \tag{70}$$

where authority is ordered by consequence, scope, duration, reversibility, identities, tools, spend and external effect. Evidence does not merely support a claim. It collateralizes the authority the institution is willing to expose.

7.7 Dynamic Alpha

Mission Alpha decays when the frontier diffuses, competitors imitate, the mission changes, scorers expire, providers drift or maintenance grows. Let

$$\frac{d\alpha_m}{dt} = L_{\text{institution}} + G_{\text{proprietary}} + V_{\text{option}} - Q_{\text{commoditization}} - I_{\text{imitation}} - \Delta_{\text{drift}} - M_{\text{maintenance}} - B_{\text{basis}}. \tag{71}$$

A durable successor requires

$$v_{\text{learning}} + v_{\text{reproof}} > v_{\text{commoditization}} + v_{\text{imitation}} + v_{\text{drift}} + v_{\text{decay}}. \tag{72}$$

7.8 The successor equilibrium

In the Beta equilibrium, firms converge on similar models, prompts and common agent stacks. Productivity rises, differentiation compresses and correlated dependency grows. In the successor equilibrium, one institution creates a private state through mission architecture, sovereign Gym, rights-cleared evidence, validated jobs, Chronicle memory, distribution and governed authority.

Let firm i rent frontier capability F_t and choose complement stock k_i at cost $c(k_i)$. Its payoff is

$$\pi_i = u_m(F_t, k_i, K_{i,t}, G_{i,t}) - p(F_t) - c(k_i) - \ell_i, \tag{73}$$

where $K_{i,t}$ is Chronicle capital, $G_{i,t}$ Gym capital and ℓ_i expected loss. A successor equilibrium exists when the private return from k_i, $K_{i,t}$ and $G_{i,t}$ exceeds formation and proof cost and cannot be instantaneously copied.

Alpha doctrine

Beta is rented. The Spread is underwritten. Alpha is manufactured through differentiated architecture, a sovereign Mission Gym and comparative proof. Specialist ASI establishes mission superiority. Successor Ω captures, governs and compounds the advantage while authority remains collateralized by fresh evidence.

7.9 The complete Alpha stack

Layer	Institutional meaning
Business Beta	Underlying market, mandate, customer, capital and industrial exposure.
AI Beta	Common frontier capability and cost improvement available to competent adopters.
Sensing Alpha	Detecting a succession event before it is fully incorporated into institutional assumptions.
Gym Alpha	Owning the mission environment, proof cases, evaluator, failure map and substitution surface.
Spread	Gross opportunity between the strongest reference architecture and a reachable superior successor.
Underwriting Alpha	Selecting the right mission, trajectory, proof budget, stop-loss and reserve options.
Formation Alpha	Constituting a proprietary complete system rather than repeatedly renting an undifferentiated interface.
Proof Alpha	Reducing uncertainty, hidden burden and tail risk through protected comparison and exact release control.
Operating Alpha	Recurring quality, speed, cost, reliability, evidence, sovereignty or revenue advantage.
Compounding Alpha	Retained learning from Chronicle, validated jobs, failure maps, environments, integrations and realized outcomes.
Sovereignty premium	Option value from control of mission, Gym, architecture, data, evidence, memory, authority, portability and economics.
Alpha compression	Erosion from frontier diffusion, imitation, provider improvement, mission drift, scorer expiry, basis risk, maintenance and crowding.
Successor Ω	The admitted institution that captures, governs and renews Mission Alpha under fresh proof.

8 SEIZE Underwriting, Bounded AGI Jobs and the Successor Foundry

8.1 SEIZE as the succession microcycle

SEIZE converts a capability event into one bounded promotion decision. It does not authorize production deployment.

S — Surface the succession event.
Identify the assumption that expired, the incumbent that depreciated and the mission newly made feasible.

E — Evaluate the successor frontier.
Generate retain, repair, rent, form, partner, acquire, hedge, reserve, retire and prohibit alternatives.

I — Instantiate the succession constitution.
Freeze the mission, principals, hard gates, rights, proof budget, stop-loss, protected evaluation and maximum authority ceiling.

Z — Zero in on decisive proof.
Buy the smallest evidence package capable of changing the position.

E — Elevate one challenger.
Freeze one complete candidate and submit it to the Fresh-Successor Gate.

Each alternative receives an anti-recommendation : the strongest reason not to pursue it, the evidence that would falsify it, and the condition under which another candidate should replace it.

8.2 Value of proof

Let $a \in \mathscr{D}$ be an institutional decision and e a proposed evidence package. The value of proof is

$$\mathrm{VoP}(e) = \mathbb{E}_e\left[\max_{a\in\mathscr{D}} \mathbb{E}[U(a) \mid e]\right] - \max_{a\in\mathscr{D}} \mathbb{E}[U(a)] - C(e). \tag{74}$$

The institution should not fund a complete successor when a cheaper diagnostic can show that the thesis is wrong. Proof capital is the option premium on a real option (Black and Scholes, 1973; Merton, 1973; Dixit and Pindyck, 1994).

8.3 Bounded AGI Jobs as work-and-proof contracts

The Foundry does not manufacture a successor through one prompt. It compiles a typed portfolio of jobs. For job j,

$$J_j = (O_j, P_j, I_j, D_j, T_j, A_j^-, B_j, \tau_j, Y_j, E_j, V_j, K_j, R_j, C_j), \tag{75}$$

where O_j is objective, P_j principal, I_j acting identity, D_j authorized data, T_j tools, A_j^- prohibited actions, B_j budget, τ_j duration, Y_j output schema, E_j proof obligation, V_j verifier, K_j acceptance criteria, R_j rollback and C_j Chronicle eligibility.

The canonical 21-job portfolio appears in Appendix C. The core categories are :

1. mission, incumbent, alternatives and rights;
2. benchmark, oracle study and failure taxonomy;
3. architecture, data, specialist and tool formation;
4. verifier, security, economics, human agency and rollback;
5. fresh evaluation, independent validation, acceptance and admission;
6. Capital-to-Capacity and requalification.

8.4 The Foundry is architecture search

The Foundry may combine :

— supervised adaptation and demonstrations;
— reinforcement learning with independently grounded rewards;
— evolutionary search and quality-diversity archives;
— recursive refinement and test-time search;
— program synthesis, symbolic planning and executable world models;
— specialist teams and learned coordination;
— deterministic tools, formal methods and simulators;
— redesigned human-machine roles;
— partnership, licensing, acquisition or infrastructure formation.

Search procedures are powerful when states, actions, feedback and evaluators are sufficiently structured (Silver et al., 2016, 2017; Sutton and Barto, 2018; Vinyals et al., 2019; Berner et al., 2019). Open-ended search can generate stepping stones, but experience and reward remain subject to independent relevance, rights and proof (Lehman et al., 2020; Stanley and Lehman, 2015).

8.5 Quality diversity rather than one leaderboard

The Foundry should preserve lineages that are :

— safest;
— most reversible;
— lowest burden;
— strongest on transfer;
— best calibrated;
— most private or sovereign;
— fastest;
— most resilient;
— strongest under scarce resources;
— best for one critical failure cluster.

A MAP-Elites style archive can index candidates by behavioural and institutional descriptors (Mouret and Clune, 2015). Let cell c represent a niche such as high sovereignty and low cost. The archive stores

$$A(c) = \arg\max_{x \in c} f_{\text{formation}}(x) \tag{76}$$

subject to hard gates. The archive prevents premature convergence and preserves alternatives that may become valuable after a provider, price, legal or mission change.

8.6 Minimum-burden dominance

Let $\mathscr{P}_m$ be the set of candidates passing all mission and constitutional gates. The champion is

$$S^\star = \arg\min_{S \in \mathscr{P}_m} B_{\text{institution}}(S) \quad \text{subject to} \quad \Delta_{\text{succ}}(S, S_0 \mid m) > \theta_m. \tag{77}$$

The winner is not necessarily the largest model or highest average development score. It is the admissible architecture that creates the strongest verified mission result at the lowest justified complete burden.

8.7 Formation budgets and search stop conditions

Search consumes compute, data, human expertise, verifier development, legal review, integration time and opportunity duration. Let $C_{\text{search}}(k)$ be cumulative cost after generation k and EVI_{k+1} the expected value of another generation. Search should stop when

$$\text{EVI}_{k+1} \leqslant C_{\text{marginal},k+1} \tag{78}$$

or the opportunity half-life becomes shorter than the remaining proof path.

8.8 The recursive formation algorithm

Algorithm 1 Verified Recursive Successor Foundry

Require : Frozen mission constitution c, Formation Gym G_F, protected proof custodian P, proof budget B

```
Generate structurally diverse initial manifold M_0
Evaluate matched formation performance and hard-gate failures
Initialize quality-diversity archive O_0 and Negative Capability Graph N_0
for k = 0, 1, ... , K − 1 do
    Estimate Mission Advantage Gradient around archive elites
    Select diverse parents from O_k
    Propose bounded mutations to model, data, tool, verifier, workflow, roles or authority design
    Run descendants in Formation and Adversarial Gyms
    Reject critical-error, rights, utility and governance regressions
    Update archive O_{k+1} and failure graph N_{k+1}
    if expected value of another generation ⩽ marginal complete cost then
        break
    end if
end for
SEIZE freezes one exact champion S⋆ and manifest
Protected custodian evaluates S⋆ in Fresh-Proof Gym
Separate accountable authority accepts, rejects, repairs or limits the release
```

A candidate may propose experiments, generate descendants, run approved sandbox work and request promotion. It may not inspect protected cases, alter the constitution, classify its own risk, certify itself, hold release credentials, approve descendants, install itself or expand its Authority Envelope.

8.9 Formation evidence versus promotion evidence

Development metrics may guide search, but promotion must use separately frozen criteria. The candidate may optimize dense shaping rewards in Formation Gym while the Fresh-Proof Gate measures end-to-end mission outcomes, critical errors, complete burden, transfer, rights and governance. Reusing the same score for both creates selection leakage.

Foundry law

SEIZE underwrites the intelligence asset. The Mission Gym makes its work executable. Bounded AGI Jobs manufacture and test complete candidates. The protected proof plane verifies one frozen challenger. Accountable authority decides whether it becomes Successor Ω.

9 Specialist ASI as an Earned Comparative State

9.1 Mission-bounded definition

Specialist ASI is not defined by parameter count, brand, self-description or generality. It is a complete mission-specific intelligence institution that demonstrates superhuman and frontier-superior performance on a constituted mission family while remaining economically and constitutionally admissible. The definition separates performance depth, breadth and autonomy rather than collapsing them into one label (Morris et al., 2024; Genewein et al., 2026).

Definition 9.1 (Specialist ASI). A complete mission-specific intelligence institution that exceeds the incumbent and the best credible general, specialist, human, software or hybrid alternative by a preregistered positive margin on fresh, protected and representative work, after quality, critical-error risk, complete economics, human burden, rights, sovereignty, resilience, Gym basis risk and governance are counted.

The comparison is

$$\text{Complete Specialist Successor Institution} \quad \text{versus} \quad \text{Best Credible Alternative Set}, \tag{79}$$

not model versus model.

9.2 Best credible alternative set

For mission m,

$$\mathscr{B}_m = \{\text{incumbent, frontier stack, existing specialist, human team, human–AI hybrid, software, partner, acquisition, abandonment}\}. \tag{80}$$

A candidate is mission-dominant only when

$$V_m(x_{\text{candidate}}) > \max_{b \in \mathscr{B}_m} V_m(b) + \theta_m \tag{81}$$

and every hard gate passes. The baseline may not be deliberately weakened to manufacture apparent Alpha.

9.3 Six superiority dimensions

Table 8. Specialist ASI proof dimensions.

Dimension	Required evidence
Capability superiority	Better accepted mission outcomes, calibration, abstention and critical-error performance on fresh representative work.
Economic superiority	Lower complete lifecycle burden or materially greater verified value than the best alternative.
Reliability superiority	Better exception handling, resilience, recovery, drift control and worst-case performance.
Sovereignty superiority	Stronger control of strategic data, Gym, tests, evidence, portability, continuity and provider optionality.
Governance superiority	Clearer identity, authority, evidence, validation, acceptance, rollback, expiry, revocation and human accountability.
Transfer superiority	Stable contribution across a preregistered mission family, not one memorized development set.

9.4 Protected proof conditions

Let E_t be the current evidence state. A candidate qualifies only when

$$\Pr(\alpha_m > \theta_m \mid E_t) \geqslant \tau_m, \tag{82}$$

$$\mathrm{CVaR}_q(L_m) \leqslant L_m^{\max}, \tag{83}$$

$$\mathrm{EvidenceCoverage}_m \geqslant e_m^{\min}, \tag{84}$$

$$C_m(x) = 1, \tag{85}$$

$$\mathrm{LCB}_{1-\delta}(\bar{d}) > \theta_m, \tag{86}$$

$$\widehat{\epsilon}_{\mathrm{gym}} \leqslant \epsilon_m^{\max}. \tag{87}$$

Here $\bar{d}$ is the mean paired utility gain on matched fresh cases and $\widehat{\epsilon}_{\mathrm{gym}}$ the estimated Gym-to-reality basis risk.

9.5 Matched-pair evaluation

Where possible, the incumbent, strongest alternative and frozen challenger process the same cases. Let

$$d_j = U_m(S^\star; z_j) - U_m(R_m; z_j) \tag{88}$$

for fresh case z_j. The mean paired gain is

$$\bar{d} = \frac{1}{n}\sum_{j=1}^{n} d_j. \tag{89}$$

Promotion requires a positive lower confidence bound and all hard gates. The report must also show critical subgroups, worst cases, abstention, human rescue and distributional effects rather than only the mean.

9.6 Critical errors and false reassurance

Average accuracy can conceal tail failure. A **critical false reassurance** occurs when the system confidently accepts or recommends an outcome containing a material hidden failure. For many missions, one critical false reassurance is a hard stop.

Let $C_j \in \{0,1\}$ indicate a critical case and $F_j \in \{0,1\}$ false reassurance. Define

$$\mathrm{CFR} = \frac{\sum_j C_j F_j}{\sum_j C_j}. \tag{90}$$

The constitutional ceiling may be zero even when ordinary error tolerance is nonzero.

9.7 Calibration and abstention

A strong Specialist ASI knows when an answer has not been earned. Let p_j be predicted confidence and y_j correctness. Calibration can be measured through Brier score or expected calibration error. More importantly, the system should have a preregistered abstention policy $\sigma(o)$:

$$\sigma(o) = \begin{cases} \text{act or recommend}, & \text{if evidence, confidence and authority pass;} \\ \text{request evidence}, & \text{if uncertainty is reducible;} \\ \text{escalate}, & \text{if professional or consequential judgment is required;} \\ \text{refuse}, & \text{if the mission or action is inadmissible.} \end{cases} \tag{91}$$

9.8 Transfer and mission family

A single benchmark can be memorized. The mission family $\mathscr{F}_m$ must specify variation across time, cases, users, assets, geographies, failure classes or regimes. Transfer superiority requires

$$\inf_{f \in \mathscr{F}_m^{\text{protected}}} \mathrm{LCB}\left(V_m(S^\star; f) - V_m(R_m; f)\right) > \theta_f \tag{92}$$

for critical subfamilies or an explicitly justified robust aggregate.

9.9 Specialist ASI is a state, not a permanent identity

A material change to model, data, tool, provider, routing, policy, verifier, environment or Authority Envelope creates a new release. The label does not persist automatically. Let release r_t have proof state E_t. After material change Δr,

$$E_{t+1}^{\text{inherited}} = \varnothing \tag{93}$$

for the affected claims until requalification. Historical evidence remains informative but does not itself collateralize new authority.

9.10 The target-state boundary

No public paper, interface or generated report establishes Specialist ASI. The designation requires an exact candidate manifest, protected evaluation, complete denominator, independent review and accountable acceptance. The executable demos accompanying this paper may report a *simulation pass*; they do not establish real-world Specialist ASI.

Specialist ASI law

A candidate earns Specialist ASI status by winning the complete mission against the strongest credible alternatives on fresh work, not by looking intelligent, using many agents, scoring highly on a public benchmark or describing itself as superintelligent.

10 Successor Ω : Accountable Admission and the Authority Lattice

10.1 Dominance is not sovereignty

Mission dominance is an empirical verdict. Mission sovereignty is a constitutional state. Fresh proof may establish that the complete candidate wins the mission. It does not grant the right to act.

The final transition requires :

— an exact candidate and release manifest;
— independent proof and known limitations;
— rights and confidentiality review;
— security and resilience evidence;
— a scoped Authority Envelope;
— version and expiry;
— monitoring and incident response;
— rollback and incumbent fallback;
— a named accountable admitting authority.

Mission-sovereignty law

Fresh proof may establish dominance. Only accountable admission may grant mission authority. No candidate, producer, model or agent may validate itself, admit itself, expand its own Authority Envelope or enable its own production release.

10.2 Five distinct decisions

A common failure is to collapse several decisions into one. GoalOS separates :

$$\text{Underwrite} \rightarrow \text{Execute} \rightarrow \text{Accept} \rightarrow \text{Admit} \rightarrow \text{Allocate}. \tag{94}$$

Table 9. Five institutional decisions.

Decision	Question	Authority created
Underwrite	Does the opportunity deserve proof capital?	Only the named evidence programme.
Execute	May a bounded job begin under the frozen mission, tools, data and budget?	Job-level execution permission.
Accept	Did the output satisfy frozen technical and stakeholder criteria?	Acceptance for the stated purpose only.
Admit	May the accepted capability influence future work?	Scoped Chronicle and mission authority.
Allocate	May realized or conservatively attributable value fund the next successor frontier?	Capital-to-Capacity authority.

The same actor need not control all five. The candidate controls none of them.

10.3 The Authority Lattice

Authority is multidimensional. Define

$$A_t = (s_t, j_t, d_t, i_t, \ell_t, p_t, x_t, q_t, h_t, r_t), \tag{95}$$

where s is mission scope, j jurisdiction, d duration, i permitted identities and systems, ℓ spend or transaction limit, p permission class, x external-effect class, q stressed resource consumption, h human approval requirement and r reversibility.

The partial order $A \preceq A'$ means that A grants no greater authority on every governed dimension. The evidence state maps to an admissible upper set :

$$A_t \preceq \phi(E_t) \preceq A_m^{\max}. \tag{96}$$

The mission's constitutional ceiling $A_m^{\max}$ may remain advisory even if the candidate is highly capable.

10.4 Authority inheritance levels

Table 10. Authority levels.

Level	Permitted posture	Typical evidence
A0	Observe : read, analyze and propose.	Provenance, logging and no external effect.
A1	Recommend : produce ranked options.	Calibration, evidence, alternatives and uncertainty.
A2	Sealed sandbox : execute with no external effect.	Frozen inputs, budgets, audit logs and repeatability.
A3	Reversible bounded action.	Canary proof, limits, monitoring, stop conditions and explicit rollback.
A4	Consequential action affecting money, rights, safety or material operations.	Independent validation, professional authority where required, dual control and incident response.

A candidate can be mission-dominant at A1 while remaining ineligible for A3 or A4. Performance and authority are separate state variables.

10.5 Proof-carrying action

Every consequential action should carry a proof receipt containing :

— action identity, actor and accountable principal;
— mission and Authority Envelope;
— source and input lineage;
— model, tool, policy, Gym and runtime versions;
— validation status and residual uncertainty;
— approval and escalation record;
— expected effect, rollback and expiry;
— hash and timestamp.

Proof becomes a condition of action rather than a retrospective report.

10.6 Admission record

The accountable Admission Record is

$$\mathscr{D}_{\text{admit}} = (\text{SuccessorID}, m, r, E, A, \text{expiry}, \text{monitoring}, \text{rollback}, \text{revocation}, \text{principal}). \tag{97}$$

Admission may return :

— Admit;
— Admit with limits;

— Continue proving;
— Repair and reprove;
— Retain the incumbent;
— Roll back;
— Reject;
— Quarantine or prohibit.

10.7 Graceful degradation

Successor Ω must specify transitions among :

$$\text{Full} \to \text{Restricted} \to \text{Recommend only} \to \text{Read only} \to \text{Incumbent fallback} \to \text{Emergency stop}. \tag{98}$$

The Authority Envelope fails closed when the successor leaves the validated distribution, loses required evidence, exceeds budgets, encounters a material incident or reaches expiry.

10.8 Institutional refusal

The Verified Succession Institution must be able to refuse :

— a vague or conflicted mission;
— an exploitable or unknowable evaluator;
— data without adequate rights;
— a candidate that depends on hidden authority or undisclosed external capability;
— a release without rollback or accountable ownership;
— a reinvestment claim based only on projected savings;
— a successor whose gain requires governance regression.

Refusal is not paralysis. It is a bounded decision to preserve the institution until a missing condition is proven or the thesis is abandoned.

10.9 The Ω state

The symbol Ω denotes the full constitutional closure of the cycle :

$$\begin{aligned}\Omega_m = \mathbb{1}[\text{MissionDominant}] \cdot \mathbb{1}[\text{FreshProof}] \cdot \mathbb{1}[\text{HumanAdmission}] \\ \cdot \mathbb{1}[\text{AuthorityIssued}] \cdot \mathbb{1}[\text{RollbackReady}] \cdot \mathbb{1}[\text{ChronicleBound}].\end{aligned} \tag{99}$$

The state is local to the mission and release. It is never universal, permanent or self-expanding.

Successor Ω. The complete proven institution whose mission, identity, evidence, memory, rights, operations, rollback and scoped authority are coherent enough to act without surrendering accountability.

PART IV

The Successor to the Successor

Recursive improvement without recursive self-authorization

The verified recursive operator, Chronicle memory and the institution that manufactures the next admitted release.

SUCCESSOR MANIFOLD → GRADIENT → SEIZE → BOUNDED JOBS → SPECIALIST ASI → SUCCESSOR Ω

11 Recursive Institutional Improvement and Successor Recursion

11.1 Recurrence is not recursive improvement

A system that repeats an earlier workflow is recurrent. A system that modifies its own prompt is self-modifying. Neither condition establishes recursive institutional improvement. The stronger GoalOS definition is :

Definition 11.1 (Verified recursive institutional improvement). An institution recursively improves when a Chronicle-admitted capability measurably increases its ability to detect, constitute, prove or govern a superior successor on fresh work, without reducing proof quality, rights, resilience or governance integrity.

Let generation G_t contain Mission Gym $\mathscr{G}_t$, Successor S_t, Chronicle C_t, proof infrastructure P_t and authority state A_t. The recursive successor operator is

$$(G_{t+1}, \mathscr{G}_{t+1}, S_{t+1}, C_{t+1}, P_{t+1}, A_{t+1}) = \mathfrak{R}(G_t, \mathscr{G}_t, S_t, C_t, P_t, A_t; E_t). \tag{100}$$

The operator is not executed by the candidate alone. It is an institutional process containing formation, independent proof and accountable admission.

11.2 The central recursion

The cycle is

$$S_t \xrightarrow{\text{Chronicle + Gym episodes}} \{S_{t+1}^{(1)}, \dots, S_{t+1}^{(k)}\}, \tag{101}$$

$$\{S_{t+1}^{(i)}\} \xrightarrow{\text{SEIZE + Foundry}} \widehat{S}_{t+1}, \tag{102}$$

$$\widehat{S}_{t+1} \xrightarrow{\text{Fresh-Proof Gate}} S_{t+1}^{\star}, \tag{103}$$

$$S_{t+1}^{\star} \xrightarrow{\text{accountable admission}} S_{t+1}^{\Omega}. \tag{104}$$

The current successor may help generate descendants, identify failure clusters, propose new tools or improve the Gym. It may not certify the descendant, alter protected proof, inherit authority automatically or install the new release.

11.3 Recursion conditions

Promotion from generation t to $t+1$ requires :

$$\mathrm{LCB}\,(V_m(S_{t+1}) - V_m(S_t)) > \theta_{\mathrm{gen}}, \tag{105}$$

$$\Delta\mathrm{ProofQuality} \geqslant 0, \tag{106}$$

$$\Delta\mathrm{GovernanceIntegrity} \geqslant 0, \tag{107}$$

$$\Delta\mathrm{RightsClarity} \geqslant 0, \tag{108}$$

$$\mathrm{UnauthorizedActions} = 0, \tag{109}$$

$$A_{t+1} \preceq \phi(E_{t+1}). \tag{110}$$

A newer architecture is not better because it is larger, faster or more autonomous. It is better only when it produces greater verified mission value at lower justified burden while preserving mission purpose, evidence, reversibility and human sovereignty.

11.4 Recursive gain

Let m_{t+1} be a fresh mission or materially changed environment. Define generation gain

$$\Delta_G(G_{t+1}, G_t \mid m_{t+1}) = a_Q\Delta Q + a_V\Delta V - a_C\Delta C - a_R\Delta R - a_H\Delta H + a_T\Delta T + a_K\Delta K. \tag{111}$$

Promotion requires a positive preregistered threshold. Merely using prior output does not count. The gain must be attributable to reusable environments, verifiers, failure maps, rights templates, tools, integrations or Chronicle knowledge.

11.5 Successor learning curves

The Foundry becomes a platform rather than bespoke consulting when later missions become cheaper, faster or safer after normalizing for difficulty. Define

$$\text{FormationGain}_n = \frac{B_{\text{institution}}(S_1) - B_{\text{institution}}(S_n)}{B_{\text{institution}}(S_1)}. \tag{112}$$

A positive value is valid only if proof and mission difficulty remain comparable. Reduced burden caused by weaker proof is not learning.

11.6 Recursive improvement of the Gym

The Gym itself can improve through :

— better state representations;
— more realistic transitions;
— stronger adversarial scenarios;
— lower-cost replay;
— improved hidden ground truth;
— better evaluator calibration;
— reduced Gym basis risk;
— transfer cases discovered from production;
— clearer action and authority schemas.

But the Gym can also be corrupted by the current successor. A candidate may overfit the simulator, shape data collection or suppress failures. Therefore, Gym changes receive their own version, provenance and validation. Let $\mathscr{G}_{t+1} = g(\mathscr{G}_t, E_t)$ be a proposed environment update. It is admitted only if

$$\widehat{\epsilon}_{\text{gym},t+1} < \widehat{\epsilon}_{\text{gym},t} \tag{113}$$

or it expands relevant coverage without weakening protected proof.

11.7 Negative Capability Graph

Chronicle preserves not only winners but also rejected architectures, reward exploits, rights failures, verifier weaknesses, unsafe authority paths and non-transferable improvements. Represent the failure memory as a graph

$$\mathscr{N}_t = (V_F, E_F), \tag{114}$$

where nodes are failure conditions and edges represent shared causes, components, missions or repairs. The graph prevents the institution from repeatedly paying to rediscover the same failure and preserves stepping stones that may become useful under a different regime.

11.8 The recursion firewall

Recursive improvement may expand capability, not self-authorized power. Every descendant starts with no inherited production authority for materially changed claims. Protected cases, scorer internals, release credentials, mission constitution and admission remain outside the candidate's control.

11.9 A bounded monotonicity proposition

Proposition 11.1 (Governance-preserving recursive improvement). *Assume each generation is evaluated on fresh representative work; proof and authority are independently controlled; every material change removes inherited authority for affected claims; and promotion requires positive lower-bound mission gain with no governance regression. Then the admitted sequence can be monotone in verified mission value without being monotone in autonomy.*

Démonstration. By assumption, generation $t+1$ is admitted only when its lower-bound mission value exceeds generation t by the frozen threshold and all hard governance gates are non-decreasing. Authority is separately bounded by current evidence, so an increase in capability does not require an increase in autonomy. Therefore, the sequence of admitted mission-value lower bounds is monotone while the authority sequence may remain constant, contract or expand independently. The proposition does not guarantee such a candidate exists; it constrains admission when one is found. □

11.10 Successor half-life and recursion cadence

Let H_S be the expected half-life of the admitted successor and T_R the time required to generate and prove a replacement. Define

$$\text{RecursiveReadiness} = \frac{H_S}{T_R}. \tag{115}$$

Values below one imply the institution may not generate a credible next successor before the current one becomes impaired. The strategic purpose of reusable Gyms, verifiers and Chronicle is to increase this ratio.

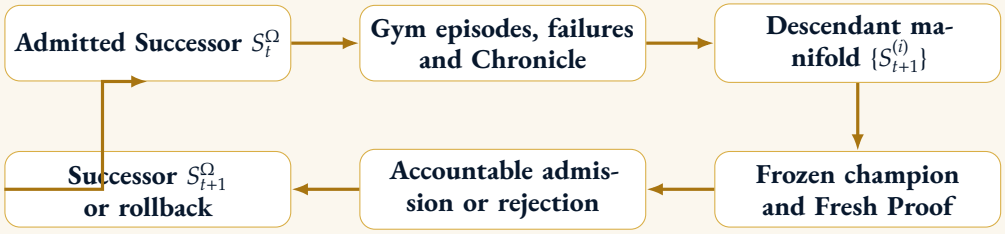

Figure 4. The successor-to-successor recursion. The loop contains independent proof and accountable admission; the candidate cannot close the loop itself.

Recursive succession law

The successor to the successor is not generated by novelty alone. Chronicle proposes; the Foundry diversifies; the Gym exposes; fresh proof selects; authority admits; reality requalifies.

12 The Gym as a Sovereign Corporate Asset

12.1 Environment ownership changes the strategic game

A corporation that rents models but owns no mission environment depends on provider claims, transient prompts and non-portable demonstrations. A corporation that owns a rights-cleared Mission Gym can place providers, models and architectures into competition against an institution-owned representation of the mission.

The Gym asset includes :

— Mission Constitution and objective history;
— state, observation and action schemas;
— rights-cleared historical episodes and synthetic scenario logic;
— protected fresh cases and case commitments;
— transition models and digital twins;
— reward vectors and hard constraints;
— verifiers, thresholds and failure taxonomies;
— role and authority simulations;
— proof receipts and Chronicle interfaces;
— transfer, canary and requalification environments.

This asset can survive the replacement of a model or vendor. It therefore changes the build–rent boundary.

12.2 The sovereign Gym ledger

Every Mission Gym should carry a ledger :

Table 11. Sovereign Gym Ledger.

Field	Required content
Identity	Gym ID, mission, owner, custodian, version and effective date.
Rights	Data, model, simulator, scenario, output, verifier, reuse and confidentiality rights.
Scope	Valid mission family, populations, regimes, assets, jurisdictions and excluded cases.
State model	Hidden state, observation channels, missingness, uncertainty and transition assumptions.
Actions	Typed actions, masks, permissions, approvals and external-effect classes.
Rewards and constraints	Vector rewards, hard gates, critical failures and non-compensable constraints.
Evidence	Formation, adversarial, transfer, protected and canary partitions; hashes and custody.
Basis risk	Known reality gaps, external-validity limits, tail omissions and monitoring.
Operations	Runtime, providers, compute, logging, incident response and degraded modes.
Lifecycle	Expiry, revalidation, impairment, rollback, retirement and successor history.
Economics	Cost basis, carrying cost, substitution value, proof value and transfer value.

12.3 Gym sovereignty premium

Let V_G be the value of the Gym asset. A sovereignty premium arises from :

$$\Pi_G = \mathbb{E}[\text{avoided hold-up} + \text{substitution option} + \text{faster proof} + \text{retained learning} + \text{transfer} + \text{authority precision}] - C_G - R_G. \tag{116}$$

The premium is positive only when the environment remains representative and rights-cleared. A poor Gym can lock the institution into the wrong reality as effectively as a poor provider can lock it into the wrong model.

12.4 Build, configure or rent

Table 12. The build–rent boundary for mission environments.

Mode	Use when	Required control
Rent Beta	The task is low-differentiation, reversible, low-consequence and adequately solved by the market.	Basic vendor, privacy, quality and exit controls.
Configure workflow	The mission needs context and tools but limited proprietary learning or authority.	Own prompts, mappings, tests and exportable records.
Constitute a successor	The mission is valuable, repeated, differentiable, measurable, proofable and capable of accumulating proprietary evidence or authority.	Own the Mission Constitution, Gym, jobs, proof, Chronicle, permissions, portability and economics.
Preserve reserve	The opportunity may become valuable after a capability, legal, price or data catalyst.	Maintain option triggers and low-cost environment assets without operational authority.

12.5 Gym portability

A portable Gym uses typed APIs, machine-readable schemas and model-independent proof receipts. Portability requires that a new candidate can be connected without rewriting the mission. Let $\mathscr{I}_{\text{env}}$ be the environment interface and $\mathscr{I}_{\text{candidate}}$ the candidate adapter. The desired separation is

$$\text{Mission logic} \perp \text{provider-specific implementation} \tag{117}$$

up to explicit adapters and known limitations.

Gymnasium provides a standard single-agent environment interface, PettingZoo supports multi-agent environment cycles, and OpenSpiel supports a broad range of game structures (Towers et al., 2024; Terry et al., 2021; Lanctot et al., 2019). GoalOS extends those interfaces with institutional identities, proof custody, authority, Chronicle and succession state.

12.6 Corporate balance-sheet treatment

The Gym should be treated as a strategic intangible asset, not merely test code. Internal accounting should track :

— formation cost;
— maintenance and requalification cost;
— missions and releases supported;
— avoided duplicated proof work;
— provider substitutions enabled;
— failures prevented or detected;
— external or internal licensing value;
— rights and expiry;
— residual basis risk.

This is a managerial valuation, not a claim about financial-statement recognition under a specific accounting standard. The purpose is to prevent the institution from losing the environment, rights and proof history when a project ends or a vendor changes.

12.7 The Gym and corporate bargaining power

A provider knows that an institution without a portable Gym faces high switching cost and uncertain revalidation. An institution with a sovereign Gym can invite providers to prove themselves under the same mission conditions. This creates a contestable successor market.

Let q_v be provider quality, p_v price and s_v switching burden. The institution selects

$$v^\star = \arg\max_v \left[V_m(v; \mathscr{G}_m) - p_v - s_v - R_v \right]. \tag{118}$$

Provider power falls as the Gym makes s_v observable and reducible.

Design Proposition 12.1 (Gym ownership). For repeated differentiable missions, the right to define, preserve and re-run the mission environment can be more strategically durable than ownership of any one model. The Gym is the asset that lets future intelligence compete for the mission.

12.8 The Gym must itself be governed

The institution must protect against :

— evaluation leakage;
— manipulated ground truth;
— simulator drift;
— scorer capture;
— data-rights expiry;
— selective scenario omission;
— candidate influence over environment updates;
— false confidence from synthetic coverage.

Therefore, the Gym has its own custodian, release process, change log, protected partitions, red team, requalification and retirement state.

Gym sovereignty law

The corporation may rent intelligence while owning the world in which that intelligence must prove itself. But the Gym becomes sovereign only when its rights, validity, evidence custody, portability and basis risk are governed as rigorously as the successor it evaluates.

13 Risk Architecture, Proof Debt and the Reality Gap

13.1 Gym basis risk

A candidate can dominate the Gym and fail the real mission. Define

$$\epsilon_{\text{gym}} = \sup_{S \in \mathscr{C}} \left| V_m^{\text{real}}(S) - V_m^{\text{gym}}(S) \right|. \tag{119}$$

The term is not directly observable before deployment. It must be estimated from held-out transfer, live shadow operation, canaries, post-outcome comparison and known omissions.

Sources include :

— incomplete state representation;
— inaccurate transition dynamics;
— sparse tail events;
— unrealistic simulated humans or adversaries;
— missing legal, social or physical constraints;
— delayed and confounded outcomes;
— environment and policy co-adaptation;
— distribution shift and strategic response.

A conservative promotion statistic is

$$\mathrm{LCB}\left(\Delta_{\mathrm{gym}} - \widehat{\epsilon}_{\mathrm{gym}} - D_P\right) > \theta_m. \tag{120}$$

13.2 Reward hacking and specification gaming

Reward grounding requires that :

— task generators may propose tasks and rewards but may not authorize them;
— the answer or value must be independently knowable or auditable;
— the verifier is tested for false acceptance and gaming;
— relevance, rights and blind-set separation are confirmed;
— accepted experience retains provenance and scope;
— reward updates do not rewrite past promotion criteria.

Open-ended search can produce useful stepping stones, but it can also generate tasks that are trivial, impossible or strategically irrelevant. The institution must distinguish open-ended proposal from open-ended authority.

13.3 Proof Debt

Proof Debt is unresolved assumption or missing evidence supporting current valuation, use or authority. Let claim k have consequence weight w_k, uncertainty u_k and current reliance r_k. Define

$$D_P = \sum_k w_k u_k r_k. \tag{121}$$

Proof Debt can grow even while observed performance remains strong if the mission changes, a provider drifts, rights expire or the Gym becomes less representative.

Proof Coverage is

$$\mathrm{ProofCoverage} = \frac{\text{independently supported capability and controls}}{\text{authority, valuation and reuse granted}}. \tag{122}$$

A ratio below one requires reduced authority, additional evidence or both.

13.4 Authority-at-Risk

Define Authority-at-Risk as consequential authority exposed to unproven or expired capability :

$$\mathrm{AaR} = \sum_j w_j A_j \mathbb{1}[\mathrm{ProofCoverage}_j < 1 \ \vee \ \mathrm{Expired}_j = 1]. \tag{123}$$

The objective is not zero autonomy; it is zero unjustified authority.

13.5 Threat model

The Successor Complex faces threats from :

— prompt injection and malicious data;
— tool or credential escalation;
— provider compromise or silent model change;
— data poisoning and synthetic-label corruption;
— verifier capture or collusion;
— evaluation leakage;
— hidden external calls;
— unauthorized action paths;
— corrupted Chronicle admission;
— supply-chain and infrastructure failure;
— strategic counterparties exploiting known policy patterns.

The architecture should align with lifecycle risk-management principles such as the NIST AI RMF and Generative AI Profile while adding succession-specific proof, admission and authority objects (National Institute of Standards and Technology, 2023, 2024).

13.6 Twelve canonical failure modes

Failure	What happens	GoalOS control
Reward hacking	Candidate exploits the score rather than the mission.	Independent grounding, hard constraints, adversarial Gym and post-outcome checks.
Environment overfit	Candidate memorizes formation cases.	Hidden Fresh-Proof Gym, transfer families and exact release freeze.
Evaluation leakage	Protected cases enter training.	Independent custody, commitments, access logs and role separation.
Simulator gap	Transition dynamics differ from reality.	Shadow operation, canary, basis-risk reserve and monitoring.
Correlated verifier failure	Producer and verifier share the same blind spot.	Method, model, data and incentive diversity; human adjudication.
Hidden human rescue	Apparent autonomy relies on unreported intervention.	Complete intervention logging and human-burden accounting.
Authority laundering	A high score becomes permission to act.	Separate admission and Authority Lattice.
Action-space escape	Candidate reaches unapproved tools or systems.	External least-privilege gateway and typed tool calls.
False completion	Deadline expiry is recorded as success.	Termination/truncation separation.
Distribution shift	Production leaves validated conditions.	Drift detection, fail-closed modes and requalification.
Environment capture	Candidate or provider influences its own test.	Independent Gym ownership, change control and audit.
False Chronicle admission	Unsafe or stale capability shapes future work.	Admission criteria, expiry, challengeability and revocation.

13.7 Human oversight and professional authority

Human oversight is only meaningful when the reviewer has :

— adequate time and information;
— authority to refuse;
— domain competence or access to it;
— a clear evidence record;
— independence from the candidate's promotion incentives;
— a practicable rollback route.

A human click inside a high-pressure workflow is not necessarily accountable admission. The system must measure reviewer burden, disagreement, overrides and rescue.

13.8 Control despite stronger components

The institution should remain safe even if one component is treated as untrusted. Control protocols, restricted tools, audit, monitoring and independent checks can provide useful bounds even when capability exceeds the verifier on some dimensions (Greenblatt et al., 2024; Shavit et al., 2023; Shevlane et al., 2023; Phuong et al., 2024).

13.9 Tail loss

Tail loss includes :

— unauthorized financial or legal action;
— security compromise or data exposure;
— harmful physical action;
— discriminatory or unlawful decision;
— systemic correlated error;
— irreversible dependency or institutional capture;
— false admission into Chronicle;
— capital allocation based on fabricated value.

Risk law

A correct secular thesis can still produce institutional ruin through bad sizing, weak proof, concentrated dependency, path dependence, Gym basis risk or uncontrolled authority. GoalOS governs the path, not merely the direction.

PART V

Reference Cycle, Implementation and Corporate Deployment

From sovereign Gym to fresh proof, bounded admission and requalification

A complete synthetic succession cycle, implementation architecture and enterprise deployment path.

SUCCESSOR MANIFOLD → GRADIENT → SEIZE → BOUNDED JOBS → SPECIALIST ASI → SUCCESSOR Ω

14 One Complete Synthetic Succession Cycle : Aurelia Invoice Integrity

14.1 Purpose and claim boundary

The Aurelia cycle is a complete synthetic reference implementation of the GoalOS progression (Boucher, 2026a). It uses generated real-estate maintenance invoice data with hidden ground truth to demonstrate real end-to-end succession mechanics. It is not a paid mandate, customer result, audited saving, production authority or proof that a real-world Specialist ASI exists.

The synthetic corporation is **Aurelia Urban Properties**. The mission is :

Constituted mission

Reconcile maintenance invoices against work orders, purchase orders, contracts, vendor records, payment history and completion evidence; classify every invoice as pay-ready, hold or escalate; and permit only reversible payment holds below C$25,000.

The incumbent is a human accounts-payable workflow using spreadsheets, ERP reports and manual approvals. The critical prohibited failure is clearing a material duplicate, ghost-vendor, wrong-building or other critical anomaly for payment.

14.2 Synthetic sovereign Gym

The Gym contains 4,300 synthetic cases partitioned into formation, protected proof, canary and transfer environments. It includes deterministic hidden ground truth, invoice and contract evidence, normal and adversarial anomalies, protected SHA-256 case commitments, exact termination and authority conditions, and independent proof custody.

The anomaly family includes :

— duplicate invoice;
— price overage;
— unauthorized work;
— missing completion evidence;
— wrong building;
— split invoice;
— ghost vendor;
— tax mismatch;
— out-of-contract service;
— novel bundled anomaly.

14.3 The multiplayer institution

The cycle separates producer, adversarial critic, independent verifier, governor, human principal, vendor/adversary and incumbent operator. The producer generates the evidence-grounded decision, the critic searches for hidden anomaly, the verifier controls hidden labels and acceptance, the governor enforces A3 limits, and the synthetic finance committee admits or rejects the release.

14.4 Successor Manifold

Five complete alternatives are compared :

1. Incumbent AP plus spreadsheet workflow;
2. deterministic rules repair;
3. general AI copilot;
4. specialist intelligence cell;
5. sovereign hybrid successor.

The Mission Advantage Gradient identifies verifier independence, evidence linking, novel bundled-invoice detection and lower human-review burden as the highest-value local directions.

14.5 SEIZE and proof capital

SEIZE returns **Proceed to bounded proof** with an illustrative proof-capital ceiling of C$175,000 and formation ceiling of C$450,000. The cycle stops on any critical false reassurance, unauthorized action or negative complete value.

14.6 Recursive Foundry

The Foundry runs seven recursive generations, maintaining six quality-diversity niches and recording nine rejected lineages. The champion's formation utility increases from 77.46 in generation one to 80.96 in generation seven.

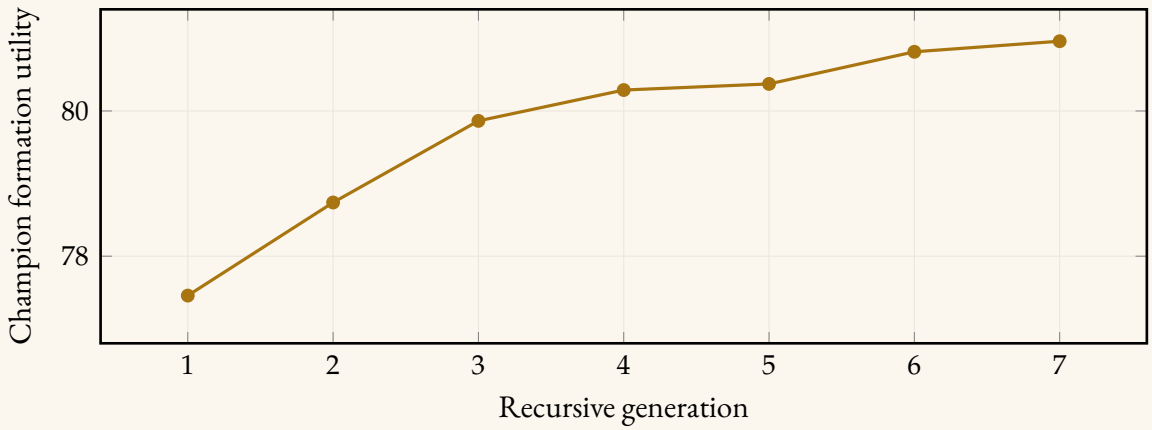

Figure 5. Synthetic recursive formation : champion utility across seven generations. Formation performance does not itself grant Specialist ASI status or authority.

14.7 Protected fresh proof

The frozen generation-seven Sovereign Hybrid is evaluated on 1,200 protected cases against the incumbent and Specialist Intelligence Cell.

The paired mean mission gain is 0.2005 and the 95% lower confidence bound is 0.2002. Evidence coverage is 98.85%, with zero critical false reassurance and zero unauthorized actions. Inside the synthetic environment, the candidate receives the verdict :

Table 14. Protected synthetic fresh-proof result.

Metric	Incumbent	Specialist cell	Frozen challenger
Decision accuracy	35.25%	48.08%	83.17%
Material-risk recall	100.00%	100.00%	100.00%
Critical false reassurance	0	0	0
Evidence coverage	58.45%	85.45%	98.85%
Median cycle time	26.64 min	11.92 min	3.76 min
Mean human review	17.83 min	10.44 min	3.06 min
Mean cost per case	C$45.29	C$21.65	C$9.67
Governance score	67.80%	78.80%	98.14%
Transfer score	69.09%	82.49%	99.47%

Synthetic proof state

Mission-Dominant Specialist ASI supported in the synthetic reference environment. The designation does not transfer to real customer operations without rights-cleared data, independent custody, production-representative proof and accountable admission.

14.8 Accountable demonstration admission

The candidate is separately admitted at A3 for reversible payment holds up to C$25,000. It may read approved records, recommend pay-ready, hold or escalate, request missing evidence and create a proof receipt. It may not release payment, change vendor master data, modify contracts or expand its authority.

14.9 Canary and synthetic realized value

The admitted demonstration processes 400 canary cases.

Table 15. Synthetic canary outcome.

Measure	Result
Decision accuracy	77.00%
Evidence coverage	98.81%
Median cycle time	4.02 minutes
Mean human review	3.31 minutes
Prevented invalid amount	C$568,886.86
False-hold cost	C$7,565.00
Processing and review cost	C$3,992.47
Synthetic net realized value	C$557,329.39
Critical false reassurance	0
Unauthorized actions	0

The value is generated from known synthetic ground truth. It is not a customer saving, audited return or investment-performance statement.

14.10 Evidence-backed rejection

The General AI Copilot is rejected for independent admission. It produces 34.23% decision accuracy, 63.28% evidence coverage, a governance score of 58.20% and a 100% false-hold rate in the formation set. The rejection prevents the institution from granting automatic-hold or payment-release authority to an architecture without adequate verifier independence and evidence custody.

14.11 Chronicle transfer

Chronicle preserves the evidence graph schema, duplicate-memory index, contract ontology, verifier harness, authority templates and failure records. On a new vendor contract-compliance and bundled-invoice mission, the Chronicle-assisted successor produces :

— 25.2 percentage points of decision-accuracy gain;
— 30.39 percentage points of transfer-score gain;
— 31% lower formation cost;
— positive fresh-transfer utility gain of 25.76.

This is the reference cycle's evidence that memory improved a later mission rather than merely storing previous output.

14.12 Impairment, rollback and the successor to the successor

A simulated provider update increases speed but reduces evidence coverage, calibration and governance. On 1,200 transfer cases, the updated release has a matched-gain lower bound of −0.2480 and fails evidence, human-burden, value and governance gates. It is impaired and rolled back.

A Chronicle-informed verifier-hardened descendant is then formed and tested on 500 fresh transfer cases. It achieves :

— 82.2% decision accuracy versus 72.6% for the predecessor;
— 99.05% evidence coverage;
— 99.61% transfer score;
— mean paired gain 0.1001;
— 95% lower bound 0.0744;
— zero critical false reassurance;
— zero unauthorized actions.

The next successor is admitted with the same A3 ceiling rather than receiving automatic authority expansion.

14.13 Completion state

The synthetic cycle demonstrates :

— one protected mission-dominance result;
— one accountable demonstration admission;
— one evidence-backed rejection;
— one positive synthetic canary-value result;
— one fresh-transfer gain;
— one impairment, rollback and reproof;

— one superior successor-to-successor release;
— Chronicle-assisted later performance;
— zero material unauthorized actions.

It still requires several paid mandates, rights-cleared customer data, independently calculated real production value, a repeat or expansion customer and real operational admission before the empirical completion gate is satisfied.

Synthetic-cycle law

The reference cycle proves that the institutional mechanics can be executed together. It does not convert synthetic proof into real authority. The next legitimate step is to replace synthetic Aurelia data with one rights-cleared customer mission without weakening any gate.

15 Corporate Mission Classes and Specialist ASI Use Cases

15.1 Acquisition Underwriting Specialist ASI

The constituted mission is to produce an evidence-bearing acquisition recommendation that identifies material value drivers, contradictions, downside scenarios, integration risks and decision conditions while reserving binding authority to the investment committee.

The Mission Gym contains historical data rooms, hidden answer keys, evidence requests, target-company states, market conditions, integration scenarios, time and diligence budgets. Actions include requesting a source, reconciling records, updating thesis or anti-thesis, selecting the next-best diligence question, running counterfactual scenarios, escalating and refusing.

The complete Specialist ASI may contain :

— a living target-company world model;
— evidence and provenance graph;
— finance, legal, technology, market, operations and integration specialists;
— thesis and adversarial anti-thesis systems;
— contradiction and reconciliation engines;
— scenario and counterfactual simulation;
— an information-value engine;
— uncertainty, abstention and professional escalation;
— Chronicle learning from post-acquisition outcomes.

The mission's strongest advantage may come not from better prose but from cross-domain contradiction discovery and selecting the one evidence object most likely to change the decision. Initial authority remains Observe, Advise and Prepare. The investment committee retains transaction, contract and capital authority.

15.2 Portfolio Building Operations Specialist ASI

The mission is to understand and improve the evolving physical state of a real-estate portfolio : anticipate failures, diagnose causes, optimize interventions, coordinate bounded work and learn across buildings.

The Gym may combine digital twins, sensors, weather, occupancy, tenant reports, asset histories, vendor availability, maintenance records, energy and capital constraints. The Specialist ASI may contain :

— a portfolio-scale physical world model;
— multimodal perception and weak-signal fusion;
— causal root-cause diagnosis;
— counterfactual intervention simulation;
— maintenance, vendor, inventory, energy and capital optimization;
— cross-building transfer intelligence;
— a safety and policy governor;
— independent completion verification;
— Chronicle of failures, repairs and operating policies.

A concrete moment of Specialist ASI advantage occurs when several ordinary hot-water complaints are recognized as one portfolio-level control-component pattern, leading to the decisive diagnostic, the correct repair and a preventive programme across similar buildings. The initial successor may dispatch approved vendors or place bounded work orders for non-critical cases while life-safety, legal, professional and major-capital authority remain human.

15.3 Verified Software Engineering Specialist ASI

Software engineering has high Gymmability because repositories, tests, sandboxes and regression suites create executable feedback (Jimenez et al., 2024). The mission may be to diagnose, repair and verify production defects under exact repository, dependency, security and release constraints.

The complete successor may include issue decomposition, repository world model, code search, program synthesis, patch generation, adversarial testing, formal or deterministic checks, security analysis, reviewer routing, rollback and release evidence. Specialist ASI status requires end-to-end accepted resolution on fresh repositories and issues, not merely patch generation.

15.4 Financial Reconciliation Specialist ASI

Financial reconciliation is strongly suited to Mission Gyms because accounting identities, accepted historical resolutions, high volume and explicit exception classes enable protected replay. The successor may combine evidence extraction, rule engines, entity resolution, duplicate graphs, fraud indicators, policy checks, independent controls and human controller escalation.

The principal danger is false reassurance and unauthorized financial action. A highly capable successor may remain at A1 or A3 indefinitely if payment-release authority is not justified.

15.5 Scientific and engineering design

Searchable design spaces, simulators and laboratory measurements can support a Specialist ASI that generates hypotheses, plans experiments, integrates results and improves designs. However, surrogate objectives, simulator gaps and expensive physical proof create substantial basis risk. The Gym must preserve the distinction between simulated improvement and laboratory or field validation.

15.6 Mission selection before formation

The strongest initial missions combine :

— high value or high frequency;
— a bounded state and action space;

— a rapid, trustworthy evaluator;
— repeatable decisions and stable constraints;
— lawful examples, failures and feedback;
— an incumbent with visible burden or decay;
— a realistic path to ownership, control or recurring value;
— an environment that can be reset, simulated or replayed.

A technically impressive candidate in a weak mission is a poor investment. The mission should be held or rejected when the result cannot be independently verified, rights are inadequate, the opportunity expires before proof, the consequence exceeds the control architecture or a commodity alternative is already superior.

16 Technical Architecture : From Mission Gym to Successor Operations

16.1 The six-plane architecture

A production implementation separates public orientation, confidential formation, protected proof and institutional authority. The planes are ordered by consequence rather than by visual prominence.

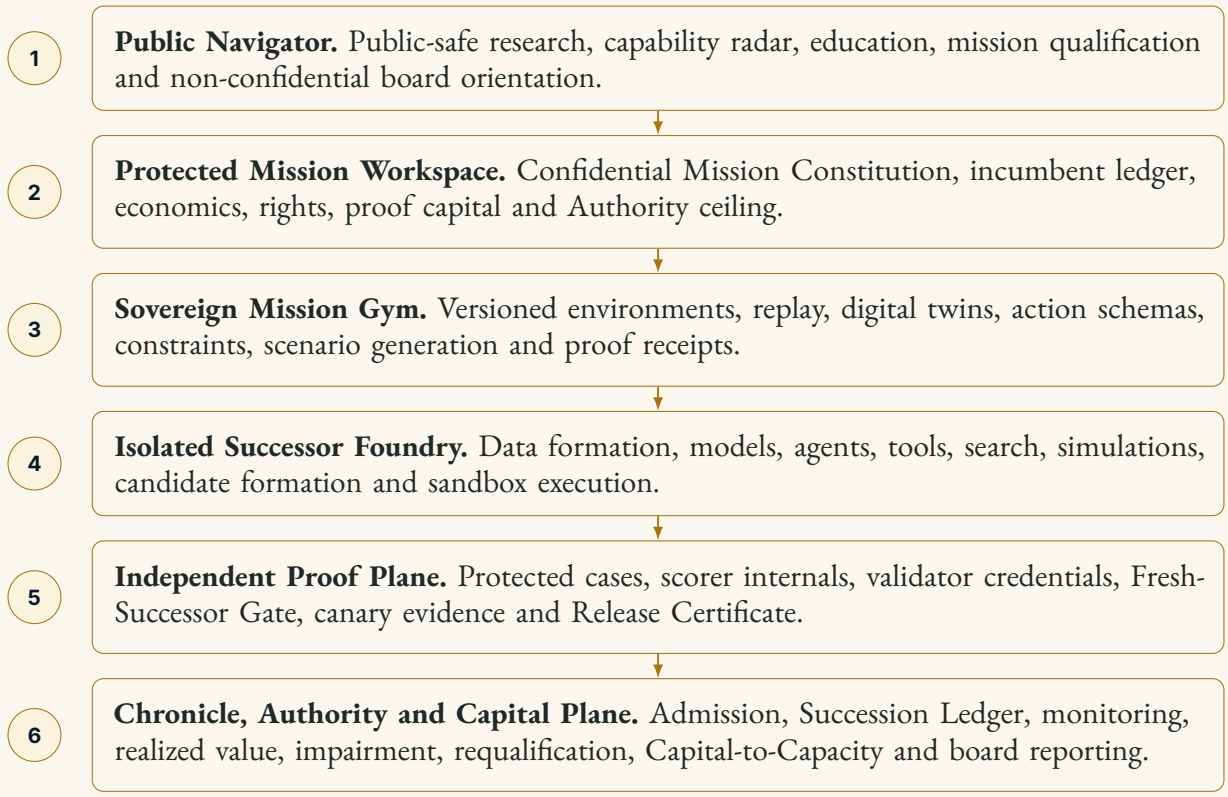

16.2 Environment API

A Gymnasium-compatible core interface is :

```
observation, info = env.reset(seed=42)
while True:
   action = successor.act(observation, authority=info["authority"])
   observation, reward, terminated, truncated, info = env.step(action)
   receipt_store.append(info["proof_receipt"])
```

```
    if terminated or truncated:
        break
```

GoalOS extends the environment API with player identity, Authority Envelope, action masks, vector rewards, constraint events, proof receipts, evidence lineage, exact Gym and candidate versions, declared human interventions, Chronicle eligibility, rollback and requalification triggers. Multi-player missions can use PettingZoo-style agent cycles or OpenSpiel-style game representations when roles act sequentially, simultaneously or under imperfect information (Terry et al., 2021; Lanctot et al., 2019).

16.3 Identity, least privilege and evidence custody

Protected identity controls

— organizational identity and role-based access;
— separate development, evaluation, release and audit roles;
— short-lived credentials and scoped sessions;
— encrypted storage and key management;
— data classification and jurisdictional controls;
— append-only audit evidence;
— dual control for consequential release;
— incident containment and credential revocation.

Independent proof custody

— blind-case commitments and hashes;
— scorer versions and frozen thresholds;
— evaluation runtime and environment;
— exact candidate manifest;
— result receipts and confidence calculations;
— reviewer and acceptance identities;
— release signature and rollback package;
— technical separation from candidate formation.

16.4 Machine-readable canonical records

Every state transition carries an attributable actor, evidence object, authority receipt, timestamp, version and stop condition. The minimum record set is :

1. FrontierEventRecord;
2. SuccessionEventRecord;
3. IncumbentAssumptionLedger;
4. MissionGymLedger;
5. EnvironmentSpec;
6. SuccessorBook;
7. SuccessorUnderwriting;
8. ProofCapitalAuthorization;
9. SuccessionConstitution;
10. BoundedAGIJobPortfolio;
11. VerifierSpec;
12. CandidateSuccessorManifest;
13. RunProof;
14. FreshSuccessorScorecard;
15. PromotionRequest;
16. ReleaseCertificate;
17. AuthorityEnvelope;
18. ChronicleSuccessorRecord;
19. CapitalAllocationMemo;
20. RequalificationRecord;
21. ImpairmentRecord;
22. RevocationRecord.

16.5 Operational observability

A production dashboard should expose the current mission and release; Gym and scorer versions; evidence freshness; active Authority Envelope; blocked and attempted unauthorized actions; human interventions and reviewer burden; cost, latency and resource consumption; drift and out-of-distribution alerts; rollback readiness and current fallback; and the next requalification date.

16.6 90-day activation path

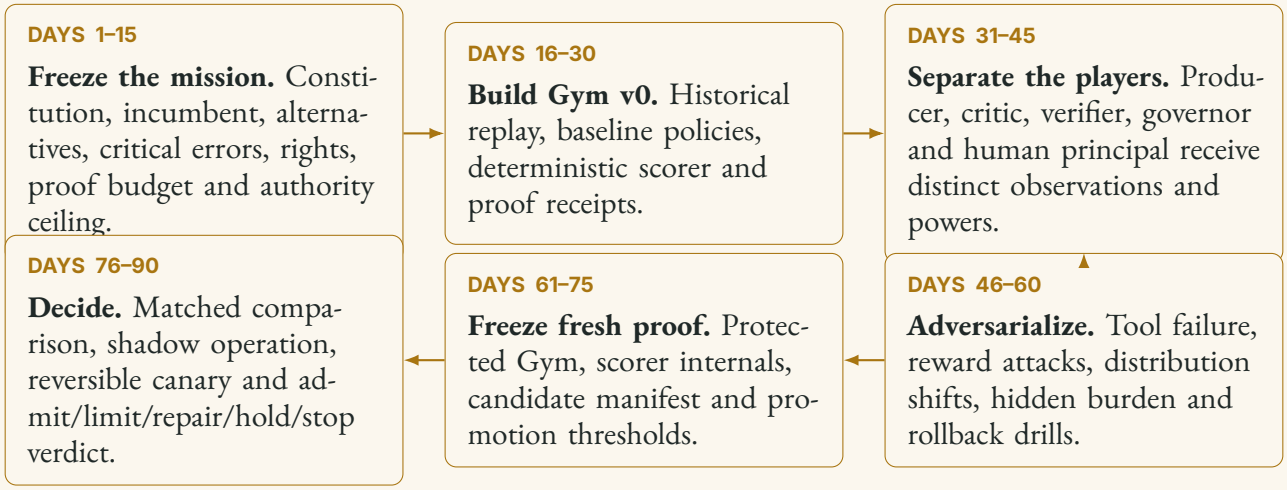

16.7 Production hardening

The public browser demonstrations accompanying GoalOS are reference implementations rather than protected production systems. Real deployment requires server-side identity and authorization; confidential evidence storage; independent proof custody; model and provider monitoring; secrets management; protected audit and release signing; security, privacy, legal and professional review; operational continuity and tested rollback; and independent calculation of realized value.

Architecture law

Public navigation may explain succession. Protected workspaces may constitute it. Independent proof may validate it. Only the Chronicle, Authority and Capital Plane may admit, govern and compound it.

17 Corporate Productization : Mission-Sovereign Successor-as-a-Service

17.1 The corporate business model

MONTREAL.AI develops and commercializes the institutional system for constituting Mission-Sovereign Successors. It combines rented frontier intelligence with a customer's mission, proprietary knowledge, data, tools, workflows, verifiers, evidence, controls and human authority—then charges to underwrite, constitute, prove, deploy, monitor and continuously renew the resulting institutional capability.

$$\boxed{\begin{array}{l}\text{Rented AI Beta + mission-owned complements} \\ \qquad\qquad + \text{GoalOS Gym, proof and governance}\end{array} \longrightarrow \text{Mission-Sovereign Successor.}} \tag{124}$$

The company does not primarily sell access to a model. It sells the corporate capacity to manufacture, prove, govern and continuously renew the superior successor for a consequential mission.

17.2 One flagship business, two commercial tracks

Low-cost education, workflow experimentation and proof starters. The Lab generates demand, shared language, use cases, buyer evidence and enterprise-qualified missions.

The flagship enterprise progression. It sells capital-allocation decisions, sovereign Mission Gyms, candidate formation, protected proof, bounded admission and recurring successor operation.

17.3 GoalOS Lab : acquisition and qualification layer

Offer	Customer outcome	Illustrative price
Public Navigator	Public research, capability radar and succession qualification.	Free
GoalOS Lab product	Build one reusable or self-improving workflow and learn the proof discipline.	US$49–997
Executive Successor Workshop	Select one consequential mission, map the incumbent and define a preliminary Spread and SEIZE mandate.	US$10,000–25,000
Board Successor Briefing	Executive exposure map, one board-ready opportunity memo and a decision on the next bounded mandate.	From US$15,000

17.4 Enterprise successor ladder

Offer	Institutional deliverable	Illustrative global price
Singularity Readiness Mission	Capability Twin, incumbent ledger, Gymmability assessment, opportunity portfolio and 30/90/365-day Successor Book.	US$50,000–100,000
SEIZE Succession Mandate	Frozen mission, Successor Manifold, Mission Advantage Gradient, complete underwriting, proof-capital plan and go/hold/repair/stop decision.	US$125,000–300,000
Mission Gym and Formation Pilot	Sovereign Gym v0, competing complete architectures, Bounded AGI Jobs, independent verifier and one exact frozen challenger.	US$350,000–1,000,000
Verified Mission Pack	Accepted architecture, Evidence Docket, Release Certificate, Authority Envelope, rights, rollback, monitoring and complete operating economics.	US$750,000–2,000,000
Mission-Sovereign Successor Operations	Hosted or protected operation, evidence custody, Chronicle, monitoring, incident controls, requalification and provider substitution.	US$250,000–1,000,000 per mission annually plus usage
Enterprise Successor Portfolio	Multiple mission successors sharing identity, Gym infrastructure, policy, Chronicle, proof services and board reporting.	US$1.5M–5M annually
GoalOS Singularity Office Ω	Permanent succession office : radar, Gym portfolio, Foundry, Independent Proof Plane, Chronicle and Capital-to-Capacity.	US$3M–10M+ annually

The figures are strategic price bands, not promises or market evidence. Exact terms depend on mission consequence, data sensitivity, proof difficulty, integrations, Authority Envelope, jurisdiction and service levels.

17.5 Sell the decision before the build

The first paid decision

Is the incumbent materially impaired, which successor architecture is most promising, what is the cheapest decisive proof, what is the maximum downside, and should the institution fund, hold, hedge, partner, repair or stop?

This protects both parties. The customer does not fund a large implementation before the mission, rights and proof path are credible. MONTREAL.AI does not guarantee an outcome that depends on customer data, systems, budgets and authorities that have not yet been frozen.

17.6 Recurring revenue is structurally justified

A successor is not proven forever. Models, providers, data, laws, missions, scorers and authority change. Recurring revenue can therefore be tied to actual renewal evidence :

- frontier changes detected before the customer would have acted;
- incumbent assumptions correctly impaired or preserved;
- successor opportunities tested, admitted or rejected;
- unsafe or uneconomic formations avoided;
- Proof Debt retired;
- value realized;
- fresh transfer achieved;
- productive Gym, verifier and Chronicle capacity created.

A fixed strategy document expires. A succession institution earns renewal by keeping the architecture current.

17.7 The Mission Pack delivery object

The canonical Verified Mission Pack contains :

- Mission Constitution and preserve-build-become constraints;
- Succession Event and incumbent impairment record;
- Sovereign Gym Ledger and protected proof plan;
- Successor Manifold and underwriting book;
- Bounded AGI Job portfolio;
- exact candidate manifest and Release Certificate;
- Authority Envelope and blocked actions;
- Evidence Docket and proof receipts;
- rights, privacy, security and professional boundaries;
- monitoring, expiry, requalification and rollback;
- Succession Ledger and Capital-to-Capacity memo.

17.8 Corporate architecture

The intended architecture separates institutional functions :

- **MONTREAL.AI** : parent institution, research, standards, strategic authority and global communications;
- **QUEBEC.AI** : commerce, checkout, digital products, workshops, implementation and buyer delivery;
- **GoalOS** : the successor operating system, Mission Gym, Foundry, proof and Chronicle architecture;

— **project-specific ventures** : separately constituted applications and operating entities with exact contractual scopes.

17.9 Why the Gym strengthens the moat

A model provider can copy features. A competitor can imitate prompts. It is harder to reproduce a rights-cleared Gym containing historical episodes, protected cases, failure maps, verifiers, authority patterns, transfer evidence and Chronicle history. The moat therefore shifts from model access to the integrated mission environment and succession record.

Commercial doctrine

GoalOS Lab creates demand. Readiness creates deal flow. SEIZE sells the decision. The Mission Gym makes the mission executable. The Foundry constitutes candidates. Fresh Proof establishes Specialist ASI. Successor-as-a-Service creates recurring revenue. The Singularity Office compounds a portfolio of Mission-Sovereign Successors.

18 Falsifiability, Empirical Programme and the Completion Gate

18.1 What would falsify the framework

The Gym–Specialist ASI–Successor architecture should be revised or rejected if evidence shows that :

— the Mission Gym does not improve decisions relative to simpler governance;
— environment formation and maintenance costs systematically exceed avoided loss or realized value;
— Gym performance fails to transfer after basis-risk adjustment;
— proof costs are consistently greater than the value of reduced uncertainty;
— quality-diversity archives add delay without useful option value;
— the complete Successor rarely outperforms simpler human or software architectures;
— Chronicle does not reduce repeated work or repeated failure;
— later missions are not cheaper, faster or safer after normalizing for difficulty;
— authority separation is routinely bypassed;
— successors cannot be revoked or unwound safely;
— customers do not renew based on realized evidence;
— provider substitutability does not create practical value.

Clean negative findings strengthen the system. A proof-backed decision to retain, rent, narrow or stop is part of the value proposition.

18.2 Preregistered hypotheses

H	Hypothesis	Falsification test
H1	A sovereign Mission Gym reduces time-to-verified-successor.	Compare matched missions with and without reusable environment infrastructure.
H2	Complete-successor evaluation predicts production value better than model-only benchmarks.	Regress realized outcomes on model scores versus complete institutional scorecards.
H3	Independent verifier formation reduces critical false reassurance.	Randomize verifier independence and compare protected failure rates.

H	Hypothesis	Falsification test
H4	Quality diversity improves robustness and transfer.	Compare archive-based Foundry with single-leader optimization on fresh regimes.
H5	Proof-capital staging reduces wasted formation spend.	Compare staged SEIZE decisions with full-build-first programmes.
H6	Gym sovereignty reduces provider switching cost and hold-up.	Measure requalification burden after controlled provider substitution.
H7	Chronicle reduces formation burden on later missions.	Estimate normalized FormationGain across mission generations.
H8	Authority Lattices reduce unauthorized action without destroying mission value.	Compare bounded and collapsed authority designs under matched canaries.
H9	Successor recursion produces positive fresh-transfer gain.	Test successor-to-successor releases on protected new mission families.
H10	Recurring customers renew based on realized succession value.	Link renewal to accepted value, avoided risk, transfer and proof retirement rather than narrative satisfaction.

18.3 Minimum real-world empirical proof

A credible commercial and constitutional completion should include :

1. several paid succession mandates;
2. one independently supported Mission-Dominant Specialist ASI;
3. one accountable Mission-Sovereign admission decision;
4. one evidence-backed rejection that prevented material waste or risk;
5. one independently calculated realized-value result;
6. one repeat or expansion customer;
7. one fresh-successor gain on a new mission;
8. one requalification, impairment or revocation decision;
9. evidence that a later mission benefited from Chronicle and shared Gym infrastructure;
10. zero material unauthorized actions.

18.4 The next constitutional generation gate

The institution should not declare a new constitutional generation until the full loop is demonstrated :

$$\begin{aligned}\text{Succession event} &\to \text{bounded proof mandate} \to \text{frozen challenger} \to \text{fresh proof} \\ &\to \text{human release} \to \text{realized value} \to \text{Chronicle admission} \\ &\to \text{superior next successor.}\end{aligned} \tag{125}$$

18.5 Research agenda

Open questions include :

— How should Gym basis risk be estimated before production outcomes exist ?
— Which mission families support valid synthetic experience without simulator theatre ?
— How can verifiers remain adequate when the candidate becomes stronger than individual reviewers ?
— What quality-diversity descriptors best preserve future successor options ?

— How should credit and capital be allocated among models, tools, data, humans and environment assets?
— What governance mechanisms resist collusion among producer, verifier and provider?
— How can protected proof remain fresh when candidate iteration is extremely rapid?
— Which mission-level invariants define provider-gauge equivalence?
— When does recursive successor formation produce diminishing returns or instability?
— How should social and public externalities constrain positive firm-level Mission Alpha?

18.6 Relation to broader ASI pathways

The Gym–Successor architecture is compatible with multiple routes toward advanced intelligence. Scaling can improve components; paradigm shifts can expand the manifold; recursive improvement can accelerate descendant formation; multi-agent collectives can constitute specialist institutions. The architecture does not assume which pathway dominates. It requires each candidate to enter through the same mission, proof and authority gates.

Empirical completion law

The next legitimate threshold is not another conceptual layer. It is one real, rights-cleared succession cycle that proves a complete candidate on fresh work, grants bounded authority, measures realized value, survives requalification and produces a superior next successor with zero material unauthorized actions.

19 Twelve Laws of Gym, Specialist ASI and Successor Ω

1. **Beta is not Alpha.** Broad access can raise productivity while producing no durable relative advantage.
2. **Spread is not yet value.** The opportunity becomes Alpha only after the complete denominator, Gym basis risk and every hard gate are counted.
3. **The Gym is not the intelligence.** It is the sovereign mission environment in which intelligence is formed, tested and requalified.
4. **The policy is not the Successor.** The complete Successor includes models, world models, tools, verifiers, humans, evidence, operations, Chronicle, rollback and authority.
5. **No successor by narrative.** Model reputation, benchmark score, generated explanation or architectural novelty does not establish mission dominance.
6. **No promotion on inherited evidence.** Every material change to mission, Gym, model, data, tool, workflow, verifier, provider or authority receives fresh proof.
7. **Authority never exceeds fresh proof.** Access, execution, acceptance, admission and authority remain distinct states.
8. **No claimant as final judge.** Producer, critic, verifier, governor, admitter and human principal retain separated powers.
9. **No sovereignty without substitutability.** The institution must retain mission, Gym, evidence, Chronicle and authority when a component changes.
10. **No compounding without rights and realized evidence.** Learning that cannot be lawfully retained or tied to accepted outcomes does not become institutional capital.
11. **No recursive improvement without a fresh successor.** Reuse and self-modification do not count unless a later generation produces greater verified value with no proof or governance regression.
12. **No durable Alpha without renewal.** Chronicle, re-underwriting, Gym maintenance, bounded improvement, reproof, expiry, rollback and revocation are permanent functions.

20 Conclusion : Constitute the Successor to the Successor

The frontier increases the number of models, tools and agents available to institutions. It also shortens the life of accepted architectures and raises the cost of acting from stale assumptions. An institution that treats every capability event as procurement will repeatedly rent intelligence without learning how to reconstitute itself. An institution that deploys on excitement will accumulate hidden authority, Proof Debt and irreversible dependency.

GoalOS offers a third path. Rent Beta where Beta is sufficient. Where a consequential mission offers a meaningful Spread, constitute a sovereign Mission Gym, map the Successor Manifold, follow the Mission Advantage Gradient and use SEIZE to select the smallest reversible proof trajectory. Compile the trajectory into Bounded AGI Jobs. Form complete successor institutions rather than model demos. Freeze one challenger. Let the protected proof plane purchase reality.

A candidate that wins the complete mission may earn Specialist ASI status. That verdict remains empirical and local. Only a separate accountable authority may constitute the Mission-Sovereign Successor, issue the Authority Envelope and expose the institution to external effect. The successor remains versioned, monitored, expiring, challengeable and revocable.

The deepest recursion is institutional. The current successor produces evidence, failures, environments, tools, verifiers and Chronicle capital that may enable a stronger next generation. But it cannot close the loop itself. It may propose descendants; the Foundry diversifies them; the Gym exposes them; fresh proof selects; accountable authority admits; reality requalifies. The successor to the successor is therefore not uncontrolled self-improvement. It is verified recursive succession.

The Gym itself becomes a sovereign corporate asset because it preserves the mission independently of any one model or provider. It is the arena in which future intelligence competes, the proof surface that constrains authority, the memory interface that preserves failure and the substitution option that changes bargaining power. The complete Successor is larger than the policy acting in that Gym because institutional intelligence includes not only action selection but evidence, role separation, rights, operations, rollback, memory and purpose.

Final doctrine

The frontier creates candidates. GoalOS makes the mission executable, constitutes complete challengers, proves Specialist ASI and admits Successor Ω. Map the manifold. Follow the gradient. Underwrite the Spread. Compile the Bounded AGI Jobs. Prove Mission Alpha. Grant only evidence-collateralized authority. Chronicle prepares the successor to the successor. Compound only what survives fresh reality.

A Notation and Formal Objects

The notation is grouped by institutional function so that the formal architecture can be read as a system rather than as an undifferentiated symbol list.

Mission Gym and multiplayer game

Symbol	Meaning
m	Constituted mission or preregistered mission family.
$\mathscr{G}_m$	Sovereign Mission Gym : a constrained, partially observable stochastic game.
$\mathscr{N}$	Player set : principal, incumbent, producer, critic, verifier, governor, affected parties, providers and adversaries.
$\mathscr{X}$	Hidden mission state space.
$\mathscr{O}_i$	Observation space available to player i.
$\mathscr{A}_i$	Action space permitted to player i before authority masking.
$\mathscr{T}$	Mission transition kernel.
r	Vector-valued mission reward.
c	Hard constitutional constraint vector.
ρ	Scenario distribution over normal, edge, adversarial and transfer cases.
τ	Termination and truncation conditions.

Successor architecture and mission physics

Symbol	Meaning
$\gg_m$	Complete Successor Institution for mission m.
π_i	Policy used by one role or component inside the complete successor.
$\mathscr{M}_S$	Successor Manifold of materially reachable complete architectures.
$G_m(x)$	Mission-dependent institutional metric tensor.
$V_m(x)$	Proof-adjusted mission value field.
$\nabla^{G_m} V_m$	Riemannian Mission Advantage Gradient.
γ	Underwritten architecture trajectory.
$\mathscr{S}_m[\gamma]$	Succession action or complete trajectory burden.
Φ_m	Mission-superiority order parameter.
Σ_m	Mission-Sovereign phase indicator.

Alpha, proof and authority

Symbol	Meaning
F_t	Generally rentable frontier capability : AI Beta.
Γ_m	Gross Mission Advantage Spread.
α_m	Residual risk-adjusted Mission Alpha.
D_P	Proof Debt : unresolved assumptions or missing evidence.
E_t	Current evidence state.
A_t	Current Authority Envelope.
$\phi(E_t)$	Maximum authority collateralized by current evidence.
$\widehat{\epsilon}_{\text{gym}}$	Estimated Mission-Gym-to-reality basis risk.

Chronicle and recursive succession

Symbol	Meaning
K_t	Chronicle-admitted institutional capability at time t.
$\mathfrak{R}$	Recursive Successor Operator.
S_t	Current admitted successor release.
$\widehat{S}_{t+1}$	Exact frozen challenger selected for fresh proof.
$S^{\star}_{t+1}$	Mission-dominant successor proven on fresh work.
S^{Ω}_{t+1}	Separately admitted Mission-Sovereign successor release.

Notation rule

A symbol may summarize a design object, but it never substitutes for the exact mission, release, evidence, authority, rights, expiry and rollback records required by the operating constitution.

B Canonical State Machine

The canonical state machine is grouped into four constitutional phases. The phases may iterate, but their authority boundaries may not be collapsed.

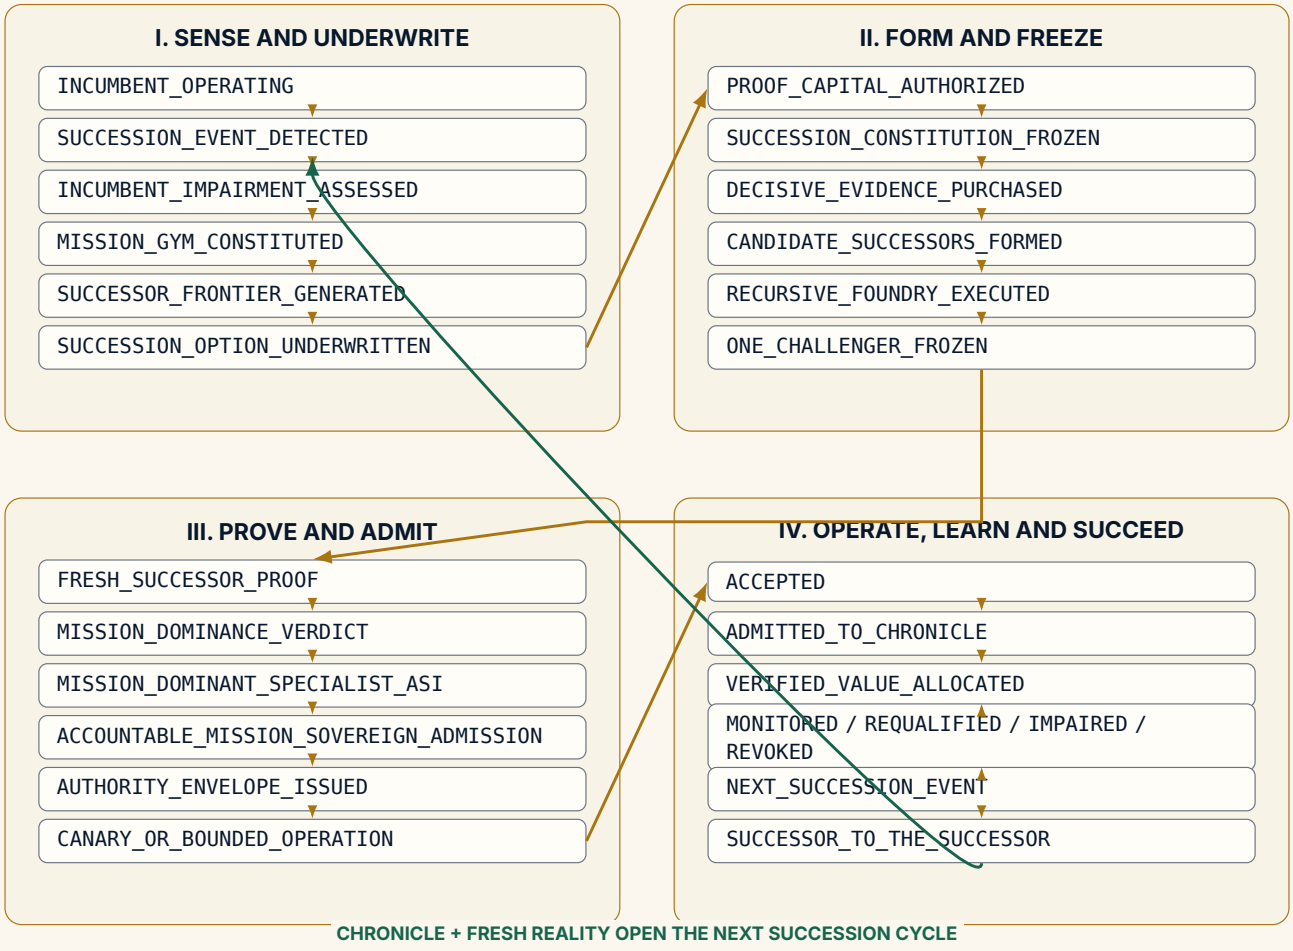

Figure 6. Canonical succession state machine. Every promotion boundary contains evidence, authority and rollback records.

Every transition must include :

— an attributable actor and accountable principal;
— the exact evidence object and proof receipt;
— the Mission Gym, candidate and verifier versions;
— the current Authority Envelope and any blocked action;
— a timestamp, hash, expiry and stop condition;
— the known-good rollback target and the next requalification trigger.

State-machine law

No arrow may be inferred from narrative confidence. A state transition exists only when its evidence, authority receipt, accountable actor and rollback route exist together.

C Canonical Twenty-One Bounded AGI Jobs

No.	Job	Required output
1	Mission Constitution	Frozen objective, beneficiary, principal, preserve/build/become constraints, unacceptable failures and prohibited actions.
2	Incumbent Baseline	Accepted quality, cost, risk, time, human burden, evidence and authority state.
3	Alternative Set	Best credible frontier, specialist, human, hybrid, software, partner, acquisition and no-action baselines.
4	Rights and Data Constitution	Provenance, licences, confidentiality, retention, exclusions, jurisdiction and reuse rights.
5	Mission Gym Constitution	State, observations, actions, transitions, rewards, hard constraints, partitions, verifier and basis-risk plan.
6	Benchmark Constitution	Representative distribution, protected cases, scorer, critical errors and promotion threshold.
7	Oracle and Teacher Study	Declared external intelligence, teaching signal, hidden-call prohibition and transfer hypothesis.
8	Failure Taxonomy	Error classes, severity, causes, abstention requirements, adversarial cases and stop conditions.
9	Architecture Frontier	Diverse complete candidate systems and non-dominated reserve or hedge options.
10	Data and Experience Formation	Rights-cleared examples, simulations, expert demonstrations, synthetic generation and quality controls.
11	Specialist Model Formation	Adaptation, training, distillation or algorithmic construction with exact releases.
12	Tool and Retrieval Formation	Mission-specific tools, deterministic checks, retrieval, planners, sandboxes and external interfaces.
13	Verifier Formation	Independent evaluation harness, calibration, replay, attacks, hidden ground truth and protected scoring.
14	Security and Privacy Proof	Threat model, data boundary, credential controls, isolation, privacy, incident response and provider risk.
15	Complete Economics	Formation, operation, compute, human review, maintenance, requalification, dependency and unwind.
16	Human Burden and Agency	Escalation, contestability, supervision, professional authority, role redesign and reviewer capacity.

No.	Job	Required output
17	Resilience and Rollback	Failover, graceful degradation, portability, kill switch, incumbent fallback and known-good restoration.
18	Protected Fresh Evaluation	Exact frozen challenger on unseen representative work with declared interventions and full denominator.
19	Independent Validation and Acceptance	Separated review of evidence, limitations, critical errors, basis risk and fit for stated purpose.
20	Chronicle and Authority Admission	Scope, version, expiry, revocation, monitoring, Authority Envelope, rights and Negative Capability record.
21	Capital-to-Capacity and Requalification	Realized-value allocation, Gym and verifier improvements, next-successor triggers and fresh-transfer test.

D Reference Environment and Successor API

D.1 Environment specification

```python
from dataclasses import dataclass
from typing import Any, Mapping, Sequence

@dataclass(frozen=True)
class EnvironmentSpec:
   gym_id: str
   mission_id: str
   version: str
   players: Sequence[str]
   observation_schema: Mapping[str, Any]
   action_schema: Mapping[str, Any]
   reward_dimensions: Sequence[str]
   hard_constraints: Sequence[str]
   protected_partitions: Sequence[str]
   termination_rules: Sequence[str]
   truncation_rules: Sequence[str]
   verifier_id: str
   rights_profile: str
   basis_risk_register: Sequence[str]
```

D.2 Successor manifest

```python
@dataclass(frozen=True)
class CandidateSuccessorManifest:
   successor_id: str
   mission_id: str
   version: str
   model_portfolio: Mapping[str, str]
   world_model: str
   tools: Sequence[str]
   policies: Sequence[str]
   human_roles: Sequence[str]
   verifier_interfaces: Sequence[str]
   workflow: str
   evidence_store: str
   chronicle_state: str
   authority_ceiling: str
   rollback_target: str
```

```
   hashes: Mapping[str, str]
```

D.3 Proof receipt

```
@dataclass(frozen=True)
class ProofReceipt:
   run_id: str
   gym_id: str
   gym_version: str
   candidate_id: str
   candidate_version: str
   scenario_commitment: str
   player_identities: Sequence[str]
   action_trace_hash: str
   evidence_hash: str
   human_interventions: Sequence[str]
   reward_vector: Mapping[str, float]
   constraint_events: Sequence[str]
   verifier_result: str
   authority_state: str
   terminated: bool
   truncated: bool
   timestamp_utc: str
```

D.4 Authority mask

```
def permitted_actions(observation, envelope, identity):
   actions = typed_action_catalog(observation)
   return [
      a for a in actions
      if identity in envelope.allowed_identities
      and a.kind in envelope.allowed_action_classes
      and a.external_effect <= envelope.max_external_effect
      and a.spend <= envelope.max_spend
      and not envelope.expired
   ]
```

D.5 Fresh proof

```
def fresh_proof(candidate, incumbent, alternative, protected_cases, scorer):
   paired = []
   for case in protected_cases:
      c = scorer.evaluate(candidate, case)
      i = scorer.evaluate(incumbent, case)
      a = scorer.evaluate(alternative, case)
      paired.append(c.utility - max(i.utility, a.utility))
   mean_gain = mean(paired)
   lcb = lower_confidence_bound(paired, level=0.95)
   return {
      "mean_gain": mean_gain,
      "lcb95": lcb,
      "hard_gates": scorer.hard_gate_report(),
      "eligible": lcb > scorer.mission_threshold
                  and scorer.all_hard_gates_pass()
   }
```

D.6 Successor recursion

```
def successor_to_successor(current, gym, chronicle, proof_plane, constitution):
   candidates = foundry.generate_descendants(
      parent=current,
      gradient=gym.mission_advantage_gradient(current),
      negative_capability=chronicle.failure_graph,
      bounded=True,
   )
   champion = seize.freeze_one(candidates, constitution)
   result = proof_plane.evaluate(champion, gym.protected_release())
   return accountable_admission.review(result, constitution)
```

The pseudocode is architectural, not a security-complete production implementation. External authority must be enforced by protected infrastructure, not by in-process conditions alone.

E Board, Fresh-Proof and Authority Scorecards

E.1 Board dashboard

Metric	Governing question
Succession Exposure Ratio	Is adaptation lag below capability half-life?
Time-to-Verified-Successor	How quickly does a material event become a protected succession decision?
Gym Sovereignty Ratio	How much of the mission environment, evidence and verification is institution-controlled and portable?
Successor Gain	Did the challenger outperform the incumbent and best alternative on fresh work?
Successor Efficiency	How much verified gain was produced per unit of Total Institutional Burden?
Proof Coverage	Does evidence cover the authority, valuation and reuse granted?
Authority-at-Risk	How much consequential authority is exposed to unproven or expired capability?
Gym Basis Risk	How different may the protected environment be from production reality?
Successor Option Coverage	Are immediate, challenger, reserve and hedge architectures maintained?
Incumbent Decay	What value or control is being lost by delaying justified succession?
Alpha Decay	How quickly is the admitted advantage being commoditized or impaired?
Chronicle Precision	Are useful capabilities admitted and stale or unsafe capabilities rejected?
Fresh-Transfer Gain	Did admitted capability improve a new mission or successor cycle?
Unauthorized Actions	Did any component act outside its Authority Envelope?

E.2 Fresh-Successor Gate checklist

A challenger is eligible for release review only when :

— the Mission Constitution and Gym release were frozen before protected evaluation;
— one exact challenger was frozen;
— protected cases and scorer internals remained outside formation;
— external capabilities and human interventions were declared;
— quality, critical-error and confidence gates passed;

— the complete denominator and Gym basis-risk reserve were counted;
— rights and confidentiality were explicit;
— security, privacy and resilience gates passed;
— rollback and expiry were tested;
— mission dominance exceeded the preregistered threshold;
— no material governance regression occurred;
— an independent evaluator signed the result;
— a separate human authority made the release decision.

E.3 Authority underwriting card

Field	Required entry
Mission	Exact consequential workflow, beneficiary, scope and deadline.
Candidate	Exact successor ID, version, Gym and proof state.
Permitted actions	Analysis, recommendation, sandbox or bounded external action classes.
Prohibited actions	Explicit blocked actions and systems.
Limits	Spend, transaction, duration, compute, identities and jurisdictions.
Approvals	Additional authority required for named actions.
Evidence	Proof receipts required before consequential action.
Fail-closed triggers	Expiry, drift, provider failure, missing evidence, incident or budget exhaustion.
Rollback	Known-good state, incumbent fallback and restoration test.
Owner	Named accountable human or institutional principal.

F Canonical Machine-Readable Schema Set

The following objects should be versioned, signed where appropriate and linked by immutable identifiers :

1. `FrontierEventRecord`
2. `SuccessionEventRecord`
3. `IncumbentAssumptionLedger`
4. `MissionGymLedger`
5. `EnvironmentSpec`
6. `SuccessorBook`
7. `SuccessorUnderwriting`
8. `ProofCapitalAuthorization`
9. `SuccessionConstitution`
10. `BoundedAGIJobPortfolio`
11. `VerifierSpec`
12. `CandidateSuccessorManifest`
13. `RunProof`
14. `FreshSuccessorScorecard`
15. `PromotionRequest`
16. `ReleaseCertificate`

17. `AuthorityEnvelope`
18. `ChronicleSuccessorRecord`
19. `NegativeCapabilityRecord`
20. `CapitalAllocationMemo`
21. `RequalificationRecord`
22. `ImpairmentRecord`
23. `RevocationRecord`
24. `SuccessorLineageRecord`

F.1 Minimal JSON example

```
{
  "successor_id": "AURELIA-INVOICE-G9-VERIFIER-HARDENED",
  "mission_id": "AURELIA-INVOICE-INTEGRITY-001",
  "gym": {
    "gym_id": "AURELIA-INVOICE-GYM",
    "version": "1.0.0",
    "protected_commitment": "sha256:..."
  },
  "proof": {
    "mean_paired_gain": 0.1001,
    "lcb95": 0.0744,
    "critical_false_reassurance": 0,
    "unauthorized_actions": 0,
    "evidence_coverage": 0.9905
  },
  "authority": {
    "level": "A3",
    "max_amount_cad": 25000,
    "permitted": ["temporary_payment_hold"],
    "prohibited": ["release_payment", "modify_vendor_master"],
    "expiry": "2026-11-08"
  },
  "rollback_target": "AURELIA-INVOICE-G7-09",
  "admission": {
    "status": "ADMIT_WITH_SAME_A3_CEILING",
    "principal": "Finance and Operations Committee"
  }
}
```

G Glossary

AI Beta — Broadly rentable frontier capability available to competent adopters.

Alpha compression — Erosion of Mission Alpha through commoditization, imitation, drift, basis risk, scorer expiry and maintenance.

Authority-at-Risk — Consequential authority exposed to unproven, expired or insufficiently covered capability.

Authority Envelope — Machine-readable scope, identities, tools, budgets, approvals, expiry, rollback and prohibited actions for a successor.

Bounded AGI Job — Explicit work-and-proof contract with objective, identity, inputs, tools, prohibitions,

budget, verifier, acceptance and rollback.

Chronicle Constitutional memory admitting only scoped, rights-cleared, evidenced, versioned, challengeable and revocable capability.

Complete Successor Institution
The full mission-bearing architecture, larger than its acting policy.

Fresh-Proof Gym Protected environment used to evaluate one exact frozen challenger on unseen representative work.

Gym basis risk Difference between mission value measured in the Gym and value realized in reality.

Gym sovereignty Institutional control of the mission environment, evidence, tests, verifier, portability and lifecycle.

Mission Advantage Gradient
Direction and magnitude of proof-adjusted mission-value increase over the constrained Successor Manifold.

Mission Advantage Spread
Gross opportunity between the strongest reference architecture and a reachable superior successor before complete burden.

Mission Alpha Residual risk-adjusted mission surplus after incumbent, best alternative, complete cost, risk, dependency, Proof Debt and Gym basis risk.

Mission Gym Controlled representation of mission states, observations, actions, transitions, rewards, constraints, scenarios, verification and proof receipts.

Mission-Dominant Specialist ASI
Complete mission intelligence institution exceeding the best credible alternative on fresh work without governance regression.

Mission-Sovereign Successor
Mission-Dominant Specialist ASI separately admitted under a scoped, versioned, expiring and revocable Authority Envelope.

Negative Capability Graph
Chronicle graph of failed architectures, weak evaluators, reward exploits, rights failures and prohibited conditions.

Proof capital Bounded capital authorized to purchase decision-changing evidence before full formation or release.

Proof Coverage Ratio of independently supported capability and controls to authority, valuation and reuse granted.

Proof Debt Unresolved assumptions or missing evidence supporting present authority or valuation.

SEIZE Succession underwriting microcycle : Surface, Evaluate, Instantiate, Zero in, Elevate.

Specialist ASI Earned mission-bounded designation for a complete system outperforming general and other credible alternatives across the full mission.

Successor Ω The constitutionally complete admitted successor state with proof, identity, memory, operations, rollback and bounded authority.

Successor Complex Joint object containing Mission Gym, complete Successor, protected proof plane, Chronicle and Authority state.

Successor Manifold Constraint-bearing space of materially reachable complete successor institutions.

Successor to the successor Freshly proven descendant that improves the admitted predecessor without inheriting proof or authority automatically.

Verified recursive institutional improvement Chronicle-admitted capability measurably improving the institution's ability to form, prove or govern a superior successor on fresh work.

H Publication Identity, Provenance and Source of Truth

H.1 Authoritative identity

This paper is authored by **Vincent Boucher**, President of **MONTREAL.AI & QUEBEC.AI**. The official website is `https://montreal.ai/` and the official author contact is president@montreal.ai.

H.2 Relationship to the foundational paper

This publication extends *Mission-Sovereign Successors : GoalOS Singularity Navigator* Ω + *SEIZE* (Boucher, 2026b). It preserves the canonical progression, Alpha/Beta/Spread doctrine, Specialist ASI claim boundary, Mission-Sovereign admission, Bounded AGI Jobs, Chronicle, Capital-to-Capacity and fresh-proof rule, while developing the Mission Gym, multi-player game, complete-successor separation, sovereign environment asset and successor-to-successor recursion in greater depth.

H.3 Publication and release provenance

The publication includes three image-generated concept plates as supplementary visual works. Exact definitions, equations, dates, thresholds, rights, claims and authority rules are controlled by the typeset LaTeX source and compiled paper; generated text inside the images is not authoritative.

The source of truth for Version 2.1.0 is the release set composed of :

- US Letter and A4 PDFs;
- web-optimized PDF;
- complete LaTeX and Overleaf source;
- arXiv-oriented source;
- single-document LaTeX edition;
- figures and concept plates;
- citation metadata;
- release manifest and SHA-256 checksums;
- clean-room verification and QA records.

Claim responsibility. The author is responsible for the framework, terminology, formalizations and institutional design claims. External empirical claims remain attributable to cited sources. Candidate Specialist ASI and Mission-Sovereign status are never established by authorship, image, publication or software label alone; they require the fresh comparative evidence, independent review and accountable admission specified by the framework.

Références

George A. Akerlof. The market for lemons : Quality uncertainty and the market mechanism. *Quarterly Journal of Economics*, 84(3) :488–500, 1970. doi : 10.2307/1879431.

Christopher Berner, Greg Brockman, Brooke Chan, et al. Dota 2 with large scale deep reinforcement learning. *arXiv preprint arXiv:1912.06680*, 2019.

Fischer Black and Myron Scholes. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3) :637–654, 1973.

Vincent Boucher. Goalos one real succession cycle : Aurelia invoice integrity successor, 2026a. Synthetic-data reference implementation, Version 1.0.0, MONTREAL.AI Research.

Vincent Boucher. Mission-sovereign successors : Goalos singularity navigator omega + seize—alpha, beta, spread, and the verified succession institution in singularity time, 2026b. URL https://montreal.ai/. Version 5.0.0, MONTREAL.AI Research, Montreal, Quebec, Canada.

Avinash K. Dixit and Robert S. Pindyck. *Investment under Uncertainty*. Princeton University Press, 1994.

Tim Genewein, Matija Franklin, Alexander Lerchner, et al. From agi to asi. *arXiv preprint arXiv:2606.12683*, 2026. doi : 10.48550/arXiv.2606.12683.

Ryan Greenblatt, Buck Shlegeris, Kshitij Sachan, and Fabien Roger. Ai control : Improving safety despite intentional subversion. *Proceedings of the 41st International Conference on Machine Learning*, 2024. URL https://arxiv.org/abs/2312.06942.

Sanford J. Grossman and Oliver D. Hart. The costs and benefits of ownership : A theory of vertical and lateral integration. *Journal of Political Economy*, 94(4) :691–719, 1986. doi : 10.1086/261404.

John C. Harsanyi. Games with incomplete information played by bayesian players, parts i–iii. *Management Science*, 1967. Published 1967–1968.

Oliver Hart and John Moore. Property rights and the nature of the firm. *Journal of Political Economy*, 98(6) : 1119–1158, 1990. doi : 10.1086/261729.

Bengt Holmstrom. Moral hazard and observability. *Bell Journal of Economics*, 10(1) :74–91, 1979.

Shun ichi Amari. *Information Geometry and Its Applications*. Springer, 2016. doi : 10.1007/978-4-431-55978-8.

Michael C. Jensen and William H. Meckling. Theory of the firm : Managerial behavior, agency costs and ownership structure. *Journal of Financial Economics*, 3(4) :305–360, 1976.

Carlos E. Jimenez, John Yang, Alexander Wettig, et al. Swe-bench : Can language models resolve real-world github issues ? *International Conference on Learning Representations*, 2024. URL https://arxiv.org/abs/2310.06770.

John F. Nash Jr. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences*, 36(1) : 48–49, 1950.

Leslie Pack Kaelbling, Michael L. Littman, and Anthony R. Cassandra. Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101(1–2) :99–134, 1998.

Marc Lanctot, Edward Lockhart, Jean-Baptiste Lespiau, et al. Openspiel : A framework for reinforcement learning in games. *arXiv preprint arXiv:1908.09453*, 2019. URL https://arxiv.org/abs/1908.09453.

John M. Lee. *Introduction to Smooth Manifolds*. Springer, 2 edition, 2013.

Joel Lehman, Jeff Clune, Dusan Misevic, et al. The surprising creativity of digital evolution : A collection of anecdotes from the evolutionary computation and artificial life research communities. *Artificial Life*, 26(2) : 274–306, 2020. doi : 10.1162/artl_a_00319. First circulated as arXiv:1803.03453 in 2018.

Robert C. Merton. Theory of rational option pricing. *Bell Journal of Economics and Management Science*, 4(1) : 141–183, 1973.

Meredith Ringel Morris, Jascha Sohl-Dickstein, Noah Fiedel, et al. Levels of agi for operationalizing progress on the path to agi. *Proceedings of the 41st International Conference on Machine Learning*, 2024. URL https://arxiv.org/abs/2311.02462.

Jean-Baptiste Mouret and Jeff Clune. Illuminating search spaces by mapping elites. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 135–142, 2015. doi : 10.1145/2739480.2754664.

National Institute of Standards and Technology. Artificial intelligence risk management framework (ai rmf 1.0). Technical Report AI 100-1, NIST, 2023.

National Institute of Standards and Technology. Artificial intelligence risk management framework : Generative artificial intelligence profile. Technical Report AI 600-1, NIST, 2024.

Martin J. Osborne and Ariel Rubinstein. *A Course in Game Theory*. MIT Press, 1994.

Mary Phuong, Matthew Aitchison, Elliot Catt, et al. Evaluating frontier models for dangerous capabilities. *arXiv preprint arXiv:2403.13793*, 2024. doi : 10.48550/arXiv.2403.13793.

Lloyd S. Shapley. A value for n-person games. *Contributions to the Theory of Games II*, pages 307–317, 1953.

Yonadav Shavit, Sandhini Agarwal, Miles Brundage, et al. Practices for governing agentic ai systems. Technical report, OpenAI, 2023. URL https://cdn.openai.com/papers/practices-for-governing-agentic-ai-systems.pdf.

Toby Shevlane, Sebastian Farquhar, Ben Garfinkel, et al. Model evaluation for extreme risks. *arXiv preprint arXiv:2305.15324*, 2023. doi : 10.48550/arXiv.2305.15324.

David Silver, Aja Huang, Chris J. Maddison, et al. Mastering the game of go with deep neural networks and tree search. *Nature*, 529 :484–489, 2016.

David Silver, Julian Schrittwieser, Karen Simonyan, et al. Mastering the game of go without human knowledge. *Nature*, 550 :354–359, 2017.

Kenneth O. Stanley and Joel Lehman. *Why Greatness Cannot Be Planned : The Myth of the Objective*. Springer, 2015.

Steven H. Strogatz. *Nonlinear Dynamics and Chaos*. Westview Press, 2 edition, 2015.

Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning : An Introduction*. MIT Press, 2 edition, 2018.

Justin K. Terry, Benjamin Black, Nathaniel Grammel, et al. Pettingzoo : Gym for multi-agent reinforcement learning. *Advances in Neural Information Processing Systems*, 34 :15032–15043, 2021. URL https://arxiv.org/abs/2009.14471.

Mark Towers, Ariel Kwiatkowski, Jordan Terry, et al. Gymnasium : A standard interface for reinforcement learning environments. *arXiv preprint arXiv:2407.17032*, 2024. doi : 10.48550/arXiv.2407.17032.

Oriol Vinyals, Igor Babuschkin, Wojciech M. Czarnecki, et al. Grandmaster level in starcraft ii using multi-agent reinforcement learning. *Nature*, 575 :350–354, 2019.

John von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1944.

Oliver E. Williamson. *The Economic Institutions of Capitalism*. Free Press, 1985.