

MONTREAL.AI RESEARCH



SOVEREIGN INTELLIGENCE INSTITUTION SERIES

GoalOS Singularity Office Ω

The Recurring Control Plane for a Living Institution

Tools to Ride the Singularity

$$\frac{dI}{dt} = \mathcal{G}(I, R, K, E, P)$$

$$\mathcal{G} = \{O, V, P, R, C\}$$

$$\frac{dG}{dt} = O - E + C$$

$$\mathcal{L} = \sum_i w_i \cdot \mathcal{U}_i$$



$$\mathcal{I}_{t+1} = \mathcal{T}(\mathcal{I}_t; \theta)$$

$$\theta^* = \arg \min_{\theta} \mathcal{L}(\theta)$$

$$\forall I \in \mathcal{I}, \exists C(I)$$



Vincent Boucher

President, MONTREAL.AI & QUEBEC.AI

Montréal, Québec, Canada \ 27 July 2026

FOUNDING OFFICE CONSTITUTION, OPERATING STANDARD,
INSTITUTIONAL ECONOMICS AND EMPIRICAL PROGRAMME

OBSERVE * GOVERN * PROVE * RECONSTITUTE * COMPOUND

Publication Record

Field	Controlled record
Title	GoalOS Singularity Office Ω
Subtitle	The Recurring Control Plane for a Living Institution
Edition	Apex Founding Office Edition — v3.0.0-SO3
Author	Vincent Boucher, President, MONTREAL.AI & QUEBEC.AI
Publisher	ENTREPRISE TERMINATOR INTELLIGENCE INC., operating as MONTREAL.AI and QUEBEC.AI
Public contact	info@quebec.ai
Publication date	27 July 2026
Document class	Founding Office Constitution, Operating Standard, Institutional Economics and Empirical Programme
Canonical doctrine	The Navigator determines where to go. The Office ensures the institution keeps deserving to go there.

Abstract

GoalOS Singularity Office Ω is a proof-governed recurring control plane for institutions operating under accelerating capability, market, regulatory and infrastructure change. It is not a one-time strategy report, model-selection dashboard or autonomous corporate actor. It maintains a live institutional state, detects material changes, reconstitutes mission architectures, assigns bounded work, requires evidence, separates execution from acceptance and memory, and admits only narrow, attributable, rights-cleared capabilities into Chronicle. The Office continuously evaluates whether the institution's current architecture remains admissible, resilient and value-creating; when it does not, the Office produces an immediate architecture, a successor architecture and a reserve option with explicit switch, stop and rollback conditions.

The paper formalizes the Office state, Institutional Pulse, Singularity Exposure Ratio, Value-of-Proof queue, architecture auction, twelve-role council, twenty-one bounded AGI Jobs, fifteen institutional receipts, Evidence Docket, encrypted Evidence Vault, Chronicle admission, fresh-successor testing, verified-value states and capital-to-capacity allocation. It also defines a six-domain monitoring constitution, recurring operating cadence, security and resilience controls, public-static and protected deployment modes, signed licensing, AGI Club access, French/English parity, privacy governance, AI transparency, Article 50 positioning, protected-deployment triggers, commercial economics and an empirical authority programme.

The central claim is institutional rather than rhetorical : when the frontier changes faster than an organization can safely revise its operating assumptions, direction is insufficient. The institution needs a recurring constitutional process that can detect what changed, prove what matters, block incomplete action, preserve human and legal authority, and compound only capabilities that continue to survive reality.

Claim Boundary

This paper specifies a testable operating institution and a production architecture. It does not claim regulatory approval, legal or tax authority, guaranteed outcomes, confidential server-side protection on a public static host, product-market fit, independently verified customer value, or demonstrated recursive self-improvement. Those claims require deployment-specific facts, professional authority and evidence generated in operation.

Executive Declaration

The world changes faster than annual planning. The Office must therefore operate as a continuous constitutional process : observe the frontier, govern authority, prove outcomes, reconstitute the institution and compound only what remains valid.



Founder's Preface : Institution Time

Institutions do not fail only because they choose the wrong destination. They also fail because the world changes while their authority, evidence, incentives, memory and productive capacity remain frozen around an earlier map.

The GoalOS Singularity Office Ω begins where the Navigator's recommendation ends. The Navigator identifies the strongest admissible direction under a defined mission. The Office makes that direction a living constitutional process. It observes the frontier and the institution, governs who may act, proves what occurred, reconstitutes what no longer fits and compounds only what remains valid.

Candidate intelligence can become abundant while trustworthy action remains scarce. Models may propose more architectures, perform more work and reduce the cost of analysis, but institutions still require accountable principals, explicit authority, attributable evidence, independent challenge, reversible action and a disciplined rule for what may enter memory.

The Founding Office Mandate

Observe without paralysis. Govern without surrendering purpose.
Prove without theatre. Reconstitute without losing continuity.
Compound without institutionalizing error.

When accepted outcomes become scoped capability, verified contribution can become capital, and capital can become compute, energy, laboratories, infrastructure and harder missions. The Office is the control plane that keeps that expansion attributable, bounded and worthy of continuation.

Vincent Boucher

President, MONTREAL.AI & QUEBEC.AI

Montréal, Québec, Canada

27 July 2026

Table des matières

Publication Record	i
Abstract	ii
Executive Declaration	iii
Founder's Preface : Institution Time	iv
I The Office Institution	1
1 The Office Thesis	2
1.1 Direction is not a durable operating state	2
1.2 The Navigator and the Office	3
1.3 The recurring control plane category	3
2 Founding Constitutional Laws	4
2.1 Twelve articles	4
2.2 Non-negotiable invariants	4
3 The Institutional State	6
3.1 State representation	6
3.2 The Institutional Pulse	6
4 Singularity Exposure and Institutional Lag	7
4.1 Singularity Exposure Ratio	7
4.2 Five concurrent clocks	7
4.3 Capability shock protocol	7
II Formal Office Model	9
5 Mission Constitution	10
5.1 Minimum constitution	10
5.2 The preserve–build–become triad	10
6 Architecture Auction and Opportunity Engine	11
6.1 Candidate architecture	11
6.2 Three-horizon portfolio	11
6.3 Recommendation stability	11
7 The Value-of-Proof Queue	12
7.1 Proof as a scarce resource	12
7.2 Proof Debt	12
8 Counterfactual Governance	13
8.1 Stress worlds	13
8.2 Institutional regret	13

9 Typed Graph Engineering	14
9.1 Five edge types	14
9.2 Why typed edges matter	14
III Proof, Authority and Memory	15
10 Authority Envelopes	16
10.1 Authority object	16
10.2 Least authority	16
11 Execute, Accept, Admit	17
11.1 Three distinct verdicts	17
12 Proof-Carrying Action	18
12.1 Action record	18
12.2 Pre-action gate	18
12.3 Post-outcome gate	18
13 Evidence Docket and Vault	19
13.1 Complete denominator	19
13.2 Cryptographic integrity	19
14 Chronicle and Selective Memory	20
14.1 Admission predicate	20
14.2 Capability record	20
14.3 Forgetting and revocation	20
15 Fresh-Successor Testing	21
15.1 Improvement function	21
IV Operating the Office	22
16 The Eighteen Command Surfaces	23
17 Twelve Office Architectures	25
18 Twelve-Role Institutional Council	26
18.1 Roles	26
19 Twenty-One Bounded AGI Jobs	27
20 Fifteen Institutional Receipts	28
21 Recurring Cadence	29
21.1 Six-domain monitoring constitution	29
V Deployment, Security and Resilience	31
22 Deployment Modes	32
23 Public Static Edition	33
23.1 Purpose	33
24 Protected Runtime	34
24.1 Architecture	34

25 AGI Club and Signed Licence Access	35
25.1 AGI Club direct ownership	35
25.2 Signed institutional licences	35
25.3 Static-gate truth boundary	35
26 Threat Model	36
27 Resilience Constitution	37
 VI Legal, Regulatory and Bilingual Operating Architecture	 38
28 Legal Positioning Doctrine	39
28.1 Legal readiness, not immunity	39
28.2 Public operating boundary	39
29 EU AI Act Article 50 Positioning	40
29.1 No categorical exemption claim	40
29.2 Application date, transitional rule and enforcement posture	40
29.3 Code-of-practice posture	40
29.4 Compliance-by-design controls	40
30 Québec Language Architecture	42
31 Privacy and Québec Law 25 Architecture	43
31.1 Data minimization	43
31.2 Privacy governance	43
31.3 Automated decisions	43
32 Contract and Publisher Architecture	44
33 Protected-Deployment Triggers	45
34 French Public Notice	46
35 English Public Notice	47
 VII Institutional Economics	 48
36 The Office Value Function	49
37 Commercial Product Architecture	50
38 AI-Native Operating Model	51
38.1 Cost discipline	51
39 Recurring Revenue and Renewal	52
40 Capital-to-Capacity Engine	53
41 The Productive-Capacity Frontier	54
 VIII Empirical Authority and Production Proof	 55
42 OfficeBench Ω	56
43 Evaluation Integrity and Secure Measurement	57
44 First Commercial Proof Run	58
45 Monitoring Metrics	59

46 1,000-Day Roadmap	60
46.1 Days 0–30 : constitute	60
46.2 Days 31–100 : prove	60
46.3 Days 101–365 : recur	60
46.4 Year 2 : productize	60
46.5 Year 3 : scale	60
47 Principal Risks and Controls	61
 IX Mission Casebook	 62
48 Enterprise Operations Office	63
49 Global Launch and Expansion Office	64
50 Finance and Treasury Office	65
51 Research Acceleration Office	66
52 Public-Interest Mission Office	67
53 Professional Intelligence Office	68
54 Portfolio Capital Office	69
55 Compute and Infrastructure Office	70
 X Constitutional Appendices	 71
A Minimum Recurring Office Contract	72
B Authority Envelope Schema	73
C Evidence Docket Schema	74
D Chronicle Capability Record	75
E Source Pack Schema	76
F Incident Register	77
G Weekly Office Brief Template	78
H Monthly Architecture Auction Agenda	79
I Article 50 Applicability Record	80
J Legal Launch Checklist	81
K Glossary	82
L AI-Generated Visual Asset Record	83
M Author and Institutional Stewardship	84
N ImageGen Institutional Visual Atlas	85
N.1 Visual provenance	90
References	91
Canonical Operating Doctrine	93

Table des figures

9.1 A simplified typed mission graph.	14
16.1 A production Office command surface from the controlled release.	24

Liste des tableaux

1.1	Canonical division of institutional labour.	3
4.1	The five clocks governed by the Office.	7
11.1	The three constitutional verdicts.	17
21.1	Office operating cadence.	29
29.1	Article 50 positioning matrix.	41
37.1	Illustrative Office product ladder; not a binding offer.	50

Première partie

The Office Institution

The Office Thesis

1.1 Direction is not a durable operating state

A navigator can identify a destination and an architecture can be correct at the moment it is accepted. Neither proposition implies that the institution will remain fit for that destination. Models change, prices collapse, laws enter into force, sources expire, employees create nexus, capital changes control, vendors fail, attack surfaces evolve and infrastructure constraints move. A living institution therefore needs a recurring control plane rather than a static recommendation.

GoalOS Singularity Office Ω is defined as the institutional process that maintains the validity of direction after direction has been chosen. It continually asks five questions :

1. What changed in the world, the frontier or the institution?
2. Which accepted claims, permissions, architectures or plans are affected?
3. What new work is required, under whose authority, with what evidence?
4. Does the current architecture remain preferred when compared with admissible successors?
5. Which verified value should become protected capability and greater productive capacity?

Office Law I — Continuing Fitness

An institution does not deserve to continue merely because its initial architecture was plausible. It deserves to continue when material assumptions remain current, authority remains valid, outcomes remain evidenced, risks remain bounded and a fresh successor has not proven superior.

1.2 The Navigator and the Office

TABLE 1.1 – Canonical division of institutional labour.

	GoalOS Singularity Navigator Ω	GoalOS Singularity Office Ω
Primary question	Where should we go?	Do we still deserve to go there?
Temporal posture	Constitution and direction	Recurrence and reconstitution
Core output	Architecture frontier and path	Current institutional state and control actions
Proof focus	Why a route is admissible	Whether the route remains valid in reality
Memory	Candidate playbooks	Chronicle-governed reusable capability
Failure response	Select another path	Pause, repair, revoke, roll back or reconstitute
Economics	Mission and launch value	Recurring verified institutional value

The two are not competitors. Navigation without an Office decays into stale strategy. An Office without navigation can preserve a locally coherent but globally inferior direction. Together they form a closed-loop institution.

1.3 The recurring control plane category

The category is not annual maintenance, generic governance, a compliance dashboard or a managed-service retainer. It is a recurring proof-and-authority institution that composes monitoring, decision architecture, bounded execution, independent evidence, acceptance, memory and reinvestment around one consequential objective.

THE RECURRING CONTROL PLANE

World and frontier signals → affected institutional claims → bounded missions → evidence → accountable decisions → Chronicle → successor architecture → productive reinvestment

Founding Constitutional Laws

2.1 Twelve articles

- Article 1. Consequential objective.** No recurring Office exists without a named objective, accountable sponsor and measurable success conditions.
- Article 2. Bounded authority.** Every person, agent, tool, provider and transaction operates under explicit, minimal, revocable authority.
- Article 3. Current evidence.** Material claims carry source, date, owner, confidence, expiry and contradiction state.
- Article 4. Separation of powers.** Intelligence, execution, proof, acceptance, memory and capital are distinct institutional functions.
- Article 5. Proof before acceptance.** A produced result cannot become an accepted outcome without satisfying its evidence and validation requirements.
- Article 6. Acceptance before memory.** Success in one context does not create a general capability right.
- Article 7. Rights before reuse.** Confidential, customer, third-party and regulated material cannot enter reusable capability without a valid rights basis.
- Article 8. Expiry and revocation.** Evidence, authority and capability may expire or be revoked when facts, sources, systems or risk change.
- Article 9. Fresh-successor evidence.** Reuse is not improvement; promotion requires a fresh comparison without governance regression.
- Article 10. Verified value.** Recommendations, discovered opportunities and modeled benefits are not realized value.
- Article 11. Reinvestment discipline.** Only verified contribution may justify productive reinvestment.
- Article 12. Mandatory-law savings.** Product language, licences and technical controls do not override applicable mandatory law or reserved professional authority.

2.2 Non-negotiable invariants

Invariant	Constitutional effect
No objective, no Office	Monitoring without a consequential objective is noise accumulation.
No principal, no authority	Every consequential decision requires an accountable natural person or institution.
No source, no material claim	Unsupported model output remains candidate intelligence.
No proof edge, no independent verdict	Producer narrative cannot attest to itself.
No acceptance, no outcome	Completion by an agent does not equal institutional success.

No rights, no Chronicle	Customer secrets and restricted material cannot become reusable platform memory.
No fresh transfer, no improvement	A successor must earn advancement on new work.
No verified contribution, no reinvestment	Booked, modeled or speculative value is not productive capital.

The Institutional State

3.1 State representation

At time t , the Office state is :

$$\mathcal{O}_t = (M_t, C_t, A_t, G_t, J_t, E_t, D_t, K_t, V_t, P_t, R_t, W_t),$$

where M is the mission constitution; C the capability twin; A authority; G the typed mission graph; J bounded jobs; E evidence; D Proof Debt; K Chronicle; V verified value; P capital and productive capacity; R risk and resilience; and W world, market, regulatory and frontier state.

The Office transition is not unrestricted optimization. It is a constrained constitutional update :

$$\mathcal{O}_{t+1} = \Phi(\mathcal{O}_t, \Delta W_t, \mathcal{A}_t, \mathcal{P}_t),$$

subject to valid authority $\mathcal{A}_t$, passed proof gates $\mathcal{P}_t$, protected rights and mandatory law.

3.2 The Institutional Pulse

The Pulse compresses the state into a board-readable operating picture while preserving drill-down. It reports :

- the Singularity Exposure Ratio;
- architecture status and recommendation stability;
- mission health and blocked dependencies;
- authority coverage and expiring permissions;
- Proof Debt by decision impact;
- evidence freshness and contradiction count;
- resilience and rollback readiness;
- accepted value, realized value and independently verified value;
- Chronicle admissions, expiries and revocations;
- the next consequential decision.

Office Law II — No False Green

An aggregate score may summarize but may not erase a fatal deficiency. A high upside score cannot compensate for missing authority, stale evidence, prohibited conduct, unbounded downside or absent rollback.

Singularity Exposure and Institutional Lag

4.1 Singularity Exposure Ratio

Let L_t denote institutional adaptation lag and H_t the half-life of mission-relevant capability assumptions. Define :

$$\text{SER}_t = \frac{L_t}{H_t}.$$

When $\text{SER} > 1$, the institution changes more slowly than its assumptions decay. When $\text{SER} < 1$, the institution may compound ahead of change, provided evidence and governance quality do not regress.

SER is not a claim that one historical threshold has been crossed. It is an operational measure of the mismatch between external change and internal adaptation.

4.2 Five concurrent clocks

TABLE 4.1 – The five clocks governed by the Office.

Clock	Question	Office response
Capability	What newly became possible, reliable or cheap?	Replicate, test and map to missions.
Evidence	What can now be established rather than asserted?	Update proof obligations and source state.
Authority	What may act, for whom and within which limits?	Issue, narrow, expire or revoke envelopes.
Market	Which customer, capital or policy assumption decayed?	Re-score architectures and trigger missions.
Capacity	Which gains deserve durable reinvestment?	Allocate only verified contribution.

4.3 Capability shock protocol

Material frontier events activate a bounded response rather than a hype cycle :

1. classify the claimed change;
2. independently reproduce it where feasible;
3. map affected assumptions, capabilities and risks;
4. auction competing execution stacks;

-
5. test under controlled mission conditions;
 6. issue an adopt, contain, monitor, preserve-against, prohibit or defer verdict;
 7. create expiry, recheck and rollback conditions.

Deuxième partie

Formal Office Model

Mission Constitution

5.1 Minimum constitution

Every Office begins from a signed or otherwise attributable Mission Constitution containing :

- consequential objective and institutional purpose ;
- accountable sponsor and delegated decision rights ;
- scope, exclusions, deadline and operating horizon ;
- permitted data, tools, providers and environments ;
- prohibited actions and sensitive-data classes ;
- acceptance criteria and independent evaluator ;
- budget, risk appetite, reversibility and stop conditions ;
- confidentiality, publication, ownership and reuse rules ;
- mandatory professional-authority dependencies ;
- Chronicle eligibility and successor-test conditions.

5.2 The preserve–build–become triad

The Mission Constitution distinguishes three objective classes :

Preserve. Protect identity, mission, rights, critical relationships, continuity, safety interlocks and the ability to reverse course.

Build. Assemble products, services, evidence systems, markets, capital, compute, laboratories, infrastructure and institutions that create measurable value.

Become. Reconstitute roles, capabilities, operating models and authority structures for a stronger frontier while preserving human agency and lawful purpose.

Architecture Auction and Opportunity Engine

6.1 Candidate architecture

A candidate Office architecture s is represented as :

$$s = (Q, C, A, G, E, R, K, V, P),$$

where Q is the operating thesis; C capability configuration; A authority model; G mission graph; E evidence plan; R resilience; K Chronicle policy; V value model; and P productive-capacity plan.

The Office does not select a route solely by expected upside. A generalized score is :

$$S(s) = w_u U + w_e E + w_a A + w_r R + w_h H + w_x X - w_c C - w_d D - w_i I,$$

where U is verified upside, E evidence strength, A authority fit, R resilience, H human agency, X compounding option value, C complexity, D Proof Debt and I irreversibility.

6.2 Three-horizon portfolio

Immediate architecture. Strongest presently admissible route.

Successor architecture. Preferred route after priority Proof Debt is retired.

Reserve option. Low-regret route preserving reversibility and optionality.

This portfolio prevents premature commitment while retaining a path to the highest-value configuration.

6.3 Recommendation stability

A recommendation is robust only when the same champion remains preferred across a declared range of assumptions. The Office reports :

- the champion's selection frequency across stress worlds;
- the margin over the runner-up;
- the most influential assumptions;
- switch and stop thresholds;
- tail outcomes and maximum modeled regret;
- the evidence most likely to change the decision.

The Value-of-Proof Queue

7.1 Proof as a scarce resource

Evidence gathering has cost, delay and opportunity value. The Office therefore prioritizes proof by expected decision impact rather than document volume.

For unresolved claim d_i , define :

$$\text{VoP}(d_i) = \frac{\Pr(\text{decision changes} \mid e_i) \cdot |\Delta V_i| \cdot M_i}{C_i + T_i + R_i},$$

where e_i is the contemplated evidence, ΔV_i the value difference between possible decisions, M_i materiality, and the denominator includes cost, time and collection risk.

Office Law III — Resolve What Can Change the Decision

The first proof job is not the easiest document to collect. It is the smallest admissible evidence package most likely to change or unlock the consequential decision.

7.2 Proof Debt

Proof Debt is not a generic confidence score. Each item records :

- exact claim or missing authority;
- affected architecture, action or value estimate;
- evidence required;
- responsible owner;
- expiry and decision deadline;
- materiality and blocking status;
- cost and delay of retirement;
- consequence if unresolved.

Counterfactual Governance

8.1 Stress worlds

The Office evaluates architectures under alternative capability, evidence, authority, capital and irreversibility conditions. A stress world is not a prediction; it is a controlled change to assumptions used to test robustness.

Core axes include :

- capability acceleration or stagnation;
- intelligence cost collapse or increase;
- evidence latency;
- provider failure or concentration;
- authority and regulatory risk;
- capital urgency and runway pressure;
- market adoption speed;
- security incident probability;
- infrastructure delay;
- human-review capacity.

8.2 Institutional regret

Let s be the chosen architecture and w a realized world. Regret is :

$$\mathcal{R}(s, w) = V(s^*(w), w) - V(s, w),$$

where $s^*(w)$ is the best admissible architecture in world w . The Office minimizes not only expected regret but also catastrophic irreversibility and governance failure.

Typed Graph Engineering

9.1 Five edge types

The Office graph distinguishes :

1. **data edges** : downstream work consumes an artifact ;
2. **proof edges** : an independent node tests a material claim ;
3. **authority edges** : work cannot advance without accountable permission ;
4. **resource edges** : jobs share or compete for budget, state, capacity or infrastructure ;
5. **governance edges** : an outcome cannot influence future missions without Chronicle admission.

9.2 Why typed edges matter

A graph that treats every arrow as a data dependency confuses computation with constitution. Some waits exist because information is unavailable. Others exist because action would be unauthorized, proof is missing, resources conflict or reuse is prohibited. Typed edges make the reason for waiting explicit and auditable.




FIGURE 9.1 – A simplified typed mission graph.

Troisième partie

Proof, Authority and Memory

Authority Envelopes

10.1 Authority object

An Authority Envelope binds an acting identity to a principal, purpose, scope and limit. At minimum it records :

- principal and delegated actor identity;
- authorized mission and purpose;
- permitted tools, systems, data classes and counterparties;
- geographic and legal perimeter;
- financial, computational and temporal limits;
- prohibited actions;
- required evidence and acceptance tests;
- escalation and human-review requirements;
- effective date, expiry and revocation method;
- rollback and incident contact.

NIST's current agent work emphasizes identity, authorization, auditing and trustworthy interoperability for software agents (NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY 2026a ; NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY 2026d). The Office extends these controls with evidence, acceptance and memory boundaries.

10.2 Least authority

Least privilege is not merely a cybersecurity setting. It is an institutional design law. The Office should grant the minimum authority required for the bounded mission and no standing authority by default.

Execute, Accept, Admit

11.1 Three distinct verdicts

TABLE 11.1 – The three constitutional verdicts.

Verdict	Question	Evidence	Effect
Execute	May the work begin?	Valid authority, input and tool controls	Bounded action becomes permitted
Accept	Did the result pass?	Tests, validation, customer or accountable review	Result becomes an accepted outcome
Admit	May the capability influence future missions?	Rights, scope, currentness, failures, transfer evidence	Narrow capability may enter Chronicle

Collapsing these verdicts creates dangerous shortcuts. Permission to act is not proof of correctness. A correct result is not a general licence to reuse its method. A reusable method is not automatically an improved successor.

Proof-Carrying Action

12.1 Action record

Every consequential action should produce a proof-carrying record :

$$PCA = (O, P, I, A, X, T, Y, E, V, C, R, K),$$

where O is objective, P principal, I actor identity, A authority, X exact input state, T tool or provider, Y result, E evidence, V validation, C acceptance, R rollback and K Chronicle disposition.

12.2 Pre-action gate

The pre-action gate passes only when :

- the objective and principal are current;
- the authority envelope is valid;
- data and tool scope are permitted;
- budget and time remain available;
- mandatory professional or human approval exists;
- rollback and incident routes are defined.

12.3 Post-outcome gate

The post-outcome gate passes only when :

- the artifact and full denominator are recorded;
- acceptance criteria are met;
- independent verification is complete where material;
- adverse outcomes and failures are disclosed;
- accountable acceptance is recorded;
- rights and reuse are determined;
- follow-up, expiry and monitoring are assigned.

Evidence Docket and Vault

13.1 Complete denominator

The Evidence Docket is not a polished executive summary. It records the complete denominator :

- mission constitution and exact facts;
- official sources and retrieval dates;
- assumptions, contradictions and sensitivities;
- models, tools, versions and prompts where appropriate;
- job inputs, outputs, retries, costs and failures;
- human and professional interventions;
- validator methods and verdicts;
- approvals, signatures and transaction receipts;
- realized and delayed outcomes;
- rights, exclusions and reuse conditions.

W3C PROV-O provides a standard vocabulary for entities, activities and agents that can support interoperable provenance (WORLD WIDE WEB CONSORTIUM 2013). The Office uses provenance as a substrate, then adds authority, acceptance and Chronicle rules.

13.2 Cryptographic integrity

The browser-local Evidence Vault may use SHA-256 for integrity, AES-256-GCM for encryption and PBKDF2-HMAC-SHA-256 for passphrase-based key derivation. These controls protect confidentiality and alteration resistance within their technical assumptions. They do not establish the truth of the underlying claim.

Integrity is not truth

A signed or encrypted record can prove that a particular byte sequence was preserved or signed. It cannot prove that an external event occurred, that the source was authoritative, that the actor had legal capacity or that the result was economically valuable. Those propositions require their own evidence and authority.

Chronicle and Selective Memory

14.1 Admission predicate

Candidate capability k may enter Chronicle only when :

$$\text{Admit}(k) = E \wedge A \wedge R \wedge S \wedge C \wedge V \wedge X,$$

where E is evidence sufficiency, A accountable acceptance, R rights clearance, S narrow scope, C currentness, V revocability and X fresh-transfer eligibility.

14.2 Capability record

A Chronicle record includes :

- exact capability statement and prohibited generalizations;
- originating missions and evidence links;
- model, tool, provider and environment versions;
- known failure conditions;
- rights and confidentiality basis;
- authorized domains and actors;
- expiry and revalidation conditions;
- revocation and supersession history;
- fresh-successor test results.

14.3 Forgetting and revocation

An institution that can only accumulate memory becomes vulnerable to stale authority, poisoned knowledge and hidden contradiction. Chronicle therefore supports reject, repair, admit-with-limits, expire, revoke and supersede decisions.

Fresh-Successor Testing

15.1 Improvement function

Let predecessor g_1 and successor g_2 perform a fresh mission m . Define :

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

where Q is quality, V verified value, T time efficiency, C cost efficiency, R risk and H human/reviewer burden.

Promotion requires a preregistered positive threshold and no material governance regression.

Office Law IV — Reuse Is Not Improvement

A successor does not earn expanded authority because it reused prior output. It earns authority only when admitted capability improves fresh work under comparable or explicitly normalized constraints.

Quatrième partie

Operating the Office

The Eighteen Command Surfaces

No.	Surface	Institutional purpose
1	Institutional Pulse	Current health, exposure, proof, authority, resilience and next decision.
2	Mission Constitution	Objective, owner, constraints, values, horizon and success conditions.
3	Apex Opportunity Engine	Immediate, successor and reserve architectures.
4	Frontier Radar	Consequential capability, regulatory, market and infrastructure changes.
5	Capability Twin	Current/target capability and preserve/own/rent/test/upgrade/s-top decisions.
6	Model & Agent Auction	Mission-specific execution-stack comparison.
7	Mission Foundry	Twenty-one bounded AGI Jobs and their dependencies.
8	Proof & Authority	Execute, Accept and Admit gates.
9	Institutional Graph	Typed data, proof, authority, resource and governance edges.
10	Counterfactual Laboratory	Robustness under alternative worlds and stress conditions.
11	Evidence Vault	Local proof bundles, encryption, restoration and purge.
12	Chronicle & Successors	Selective memory and fresh-mission improvement.
13	Capital-to-Capacity	Verified-value allocation into productive capacity.
14	Office Cadence	Continuous, weekly, monthly, quarterly and event-driven work.
15	Source Register	Signed source provenance, official status and currentness.
16	Research & MasterClass	Constitutional and operational education.
17	Legal Center	Bilingual terms, privacy, transparency, rights and deployment boundary.
18	Diagnostics	Release hashes, signature state, source packs, sessions and system health.

Ω

GoalOS Singularity Office Ω

THE RECURRING CONTROL PLANE FOR A LIVING INSTITUTION

OfficeConstitutionFrontier RadarMissionsCapability TwinArchitecture AuctionProof & AuthorityEvidenceFR ⓘ ⌵ Preview

ALGORITHMIC INSTITUTION

Scenarios, scores, agent deliberations and recommendations are generated by software. They are research and mission-design outputs—not professional advice, approval, authority, funding or proof that an external outcome occurred.

APEX BOARD ROOM Ω

Apex Board Room

Turn the complete institutional state into a decision-ready executive brief, weekly command memo and print-ready board pack.

Generate board briefDownload HTMLDownload weekly memoPrint / Save as PDF

MONTREAL.AI · v1.0.0-SO1

GoalOS Singularity Office Ω

MONTREAL.AI — Demonstration Office · 2026-07-27

Institutional verdict: ADOPT

Global Launch & Expansion Office — Continuously optimizes jurisdictions, capital, talent, public support and market entry.

SEE	PROOF POSTURE	AUTHORITY COVERAGE	MISSION HEALTH	VERIFIED VALUE
1.05	40%	82%	37%	US\$325,000

72-hour / 30 / 90 / 365-day command sequence

0–72h

Revalidate decision-changing signals, freeze the authority envelope and launch the first bounded Global Launch & Expansion Office pilot.

30 days

Demonstration loaded
Data is synthetic.

FIGURE 16.1 – A production Office command surface from the controlled release.

Twelve Office Architectures

No.	Architecture	Design logic
1	Proof-First Adaptive Office	Retires the highest-decision-impact Proof Debt before broad expansion.
2	Frontier Capture Office	Moves rapidly when a replicated frontier capability creates a short-lived option.
3	Sovereign Capability Core	Concentrates strategic data, authority and reusable capability in a controlled core.
4	Resilience-First Office	Optimizes provider optionality, rollback, continuity and low tail risk.
5	Human Agency Transition	Redesigns roles and authority while preserving human purpose and contestability.
6	Productive Capacity Office	Converts verified value into compute, laboratories, energy or infrastructure.
7	Global Launch & Expansion Office	Reconstitutes jurisdiction, capital, talent, market and infrastructure architecture.
8	High-Frequency Operations Office	Targets recurrent operational missions with measurable service outcomes.
9	Research Acceleration Office	Governs hypothesis, experiment, replication and research transfer.
10	Portfolio Capital Office	Allocates capital across a portfolio of verified mission opportunities.
11	Public-Interest Mission Office	Prioritizes transparency, public purpose, equity, impact and accountable authority.
12	Minimal-Regret Reserve	Preserves option value and reversibility under unresolved uncertainty.

Twelve-Role Institutional Council

18.1 Roles

1. **Frontier Sentinel** detects material changes and source decay.
2. **Capability Cartographer** maps impact on the Capability Twin.
3. **Model Auctioneer** compares execution stacks under normalized conditions.
4. **Mission Compiler** decomposes the objective into bounded work.
5. **Authority Governor** verifies identity, permissions and escalation.
6. **Adversarial Auditor** searches for failures, hidden assumptions and manipulation.
7. **Evidence Verifier** tests artifacts against acceptance criteria.
8. **Successor Architect** constructs fresh successor configurations.
9. **Capacity Strategist** converts verified contribution into productive capacity.
10. **Opportunity Portfolio Architect** balances immediate, successor and reserve options.
11. **Rights and Privacy Steward** governs data, confidentiality, ownership and reuse.
12. **Institutional Externalities Auditor** evaluates effects on people, markets, infrastructure and public interests.

These are constitutional roles, not claims that twelve independent frontier models are always operating. In the public-static edition they may be deterministic analytic functions. In a protected deployment they may be assigned to separate models, people or institutions, provided independence and authority are preserved.

Twenty-One Bounded AGI Jobs

No.	Job	Primary proof object
1	Constitute mission	Signed Mission Constitution
2	Map current institutional state	Capability Twin snapshot
3	Detect frontier change	Frontier Event Record
4	Validate official sources	Source Pack and retrieval receipts
5	Map affected assumptions	Impact graph
6	Generate architecture frontier	Candidate architecture register
7	Run counterfactual stress worlds	Scenario laboratory record
8	Auction models and agents	Normalized execution comparison
9	Compile bounded mission graph	Typed dependency graph
10	Issue authority envelopes	Signed/approved authority records
11	Run sandboxed work	Execution logs and artifacts
12	Collect evidence	Evidence Docket
13	Perform independent validation	Validator report
14	Obtain accountable acceptance	Acceptance receipt
15	Assess rights and reuse	Rights register
16	Determine Chronicle disposition	Admission decision
17	Construct successor	Successor architecture
18	Run fresh-successor test	Matched comparison
19	Calculate verified value	Value ledger
20	Allocate productive reinvestment	Capital-to-capacity receipt
21	Publish board and stakeholder brief	Signed executive record

Fifteen Institutional Receipts

1. Consequential objective
2. Accountable principal
3. Current official sources
4. Authority Envelope
5. Evidence Docket
6. Independent validation
7. Accountable acceptance
8. Rights and scope
9. AI-output transparency
10. Rollback and continuity
11. Currentness and expiry
12. Externalities and impact
13. Fresh-successor test
14. Chronicle admission
15. Productive-reinvestment authority

A receipt is not a decorative certificate. It is a controlled record containing issuer, scope, effective time, evidence references, expiry, revocation and the exact institutional effect.

Recurring Cadence

TABLE 21.1 – Office operating cadence.

Cadence	Core work	Decision output
Continuous	Source expiry, authority, mission, resilience and incident monitoring	Alerts, blocks and bounded response
Weekly	Frontier and Institution Brief; Proof Debt review	Priority missions and escalations
Monthly	Architecture Auction; Capability Twin and verified-value committee	Retain, repair, expand, defer or stop
Quarterly	Full reconstitution and counterfactual stress test	Successor and reserve architectures
Material event	Immediate rerun after capability, law, security, capital or ownership change	Event-specific verdict and rollback
Annual	Authority, privacy, security, provider, Chronicle and reinvestment review	Renewal constitution

NIST’s post-deployment monitoring work identifies functionality, operations, human factors, security, compliance and large-scale impacts as distinct monitoring concerns (NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY 2026c). The Office maps each concern to named metrics, owners, thresholds and evidence.

21.1 Six-domain monitoring constitution

The Office adopts six distinct monitoring domains : functionality, operations, human factors, security, compliance and large-scale impacts (RAO et al. 2026). Each requires named metrics, evidence sources, thresholds, accountable owners and escalation.

Domain	Question	Representative Office controls
Functionality	Does the system continue to do the intended work?	Mission success, regressions, capability drift, false recommendation rate
Operations	Does service remain available and controlled?	Reliability, latency, cost, capacity, dependency health, provider fallback
Human factors	Can people understand, contest and govern the system?	Review burden, override quality, training, role clarity, agency and accessibility

Security		Are identity, tools, data and actions protected?	Authorization, sandboxing, incident signals, credential boundaries, rollback readiness
Compliance		Are obligations, policies and evidence current?	Source expiry, role classification, Article 50 record, privacy controls, professional gates
Large-scale impacts	im-	What effects occur beyond the mission?	Labour, market, infrastructure, environmental, concentration and public-interest effects

Office Law XIII — Monitoring Must Change Decisions

A metric belongs in the Office only if a defined threshold can change a mission, authority envelope, architecture, Chronicle record or productive-capacity allocation. Monitoring without institutional consequence is theatre.

Cinquième partie

Deployment, Security and Resilience

Deployment Modes

Level	Mode	Boundary
1	Public reference institution	Public-safe research, simulation, education and architecture demonstration.
2	Private advisory workspace	Confidential mission design without persistent external execution.
3	Assisted execution	Approved connectors and human-in-the-loop actions under customer authority.
4	Governed transaction	Material external action requiring complete receipts, roll-back and professional authority.
5	Strategic infrastructure	Continuous monitored operation involving critical systems, physical assets or major capital.

Each higher level requires stronger identity, contracts, security, evidence, insurance, incident response and professional review. A public GitHub Pages deployment must not be represented as a confidential private runtime.

Public Static Edition

23.1 Purpose

The public-static edition provides :

- public-safe architecture and simulation ;
- browser-local mission state ;
- conditional AGI Club or signed-licence access ;
- transparent scoring and deterministic role simulation ;
- local evidence and export controls ;
- research and MasterClass materials.

It does not provide confidential server-side DRM, hidden source code, external transaction authority or guaranteed control against a technically sophisticated user modifying client-side code.

Protected Runtime

24.1 Architecture

The optional protected runtime may support live model deliberation, private source packs and approved enterprise connectors. It should be deployed separately with :

- TLS and exact origin allowlisting;
- secret and key management;
- authenticated principals and roles;
- payload limits and rate limiting;
- isolation and sandboxing;
- structured logs, retention and deletion;
- incident response and business continuity;
- no implicit external transaction authority.

Controlled sandboxes, checkpointing and durable execution can support long-running agent work (OPENAI 2026); GoalOS adds mission constitution, authority, evidence, acceptance and Chronicle.

AGI Club and Signed Licence Access

25.1 AGI Club direct ownership

Ordinary owner access may require current direct ownership of an exact one-label name :

`label.club.agi.eth`

The gate should distinguish direct Registry ownership from recognized wrapped ownership and should reject resolver control, operator approval, delegated management, expired ownership and nested names. ENS publishes official deployment information and ownership semantics (ETHEREUM NAME SERVICE 2026).

25.2 Signed institutional licences

The publisher may issue signed licences for Professional, Practitioner, Team, Institution, Protected Deployment and Sovereign tiers. A licence should bind product, compatible release, effective/expiry dates, origin, tier, seats and legal-bundle digest.

25.3 Static-gate truth boundary

A public client-side gate is a membership, provenance, licensing and conditional-use mechanism. It is not a promise of confidential DRM. Protected information belongs in a protected deployment.

Threat Model

Threat	Failure mode	Primary controls
Objective injection	Mission is silently redirected	Frozen constitution, principal approval, diff and alert
Prompt/tool injection	External content manipulates action	Isolation, input validation, tool policies, least privilege
Authority laundering	Agent exceeds or transfers rights	Signed envelope, no delegation by default, expiry/revocation
Credential exfiltration	Secrets leave protected scope	Vaulting, compartmentalization, egress controls, redaction
Evidence fabrication	Synthetic artifacts presented as reality	Independent source retrieval, provenance, validator separation
Memory poisoning	False capability enters Chronicle	Admission predicate, rights, scope, fresh transfer, revocation
Validator capture	Producer controls evaluator	Preselected independent validator, blind comparison, rotation
Provider monoculture	Correlated model or platform failure	Multi-provider auction, local fallback, portable state
Runaway action/cost	Work exceeds bounds	Budgets, duration, circuit breakers, approvals, rollback
Insider abuse	Authorized person misuses authority	Separation of duties, logs, dual control, review
Critical-system harm	Digital action causes physical or public harm	Protected deployment, professional authority, simulation, stop gates

The Office should map controls to NIST AI RMF governance, mapping, measurement and management functions (NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY 2023), and test known agent and LLM failure classes, including those summarized by OWASP (OWASP FOUNDATION 2025).

Resilience Constitution

Every critical capability should have :

- primary execution path;
- independent verification path;
- fallback execution path;
- human escalation path;
- stop condition;
- recovery procedure;
- data and state portability;
- provider substitution test;
- last-known-good checkpoint;
- post-incident evidence plan.

Maximum autonomy is not the objective. The objective is maximum useful autonomy that remains governable, observable, reversible and aligned with institutional purpose.

Sixième partie

Legal, Regulatory and Bilingual Operating Architecture

Legal Positioning Doctrine

28.1 Legal readiness, not immunity

The Office is designed to reduce unnecessary legal and regulatory exposure through narrow conduct, explicit role boundaries, transparency, data minimization, professional escalation and fail-closed controls. It does not claim exemption, immunity or approval. Applicability depends on actual facts, functionality, jurisdiction, users and deployment roles.

Do not design around slogans of exemption. Design so that the product's actual conduct is narrow, transparent, attributable and capable of satisfying the obligations that genuinely apply.

28.2 Public operating boundary

The public release is characterized as an experimental browser-local research, simulation, mission-design, evidence-governance and institutional-monitoring system. It does not itself perform legal, tax, investment, medical or other reserved professional services; execute external transactions; make high-impact decisions about individuals; file with authorities; hold customer funds; or operate critical infrastructure.

EU AI Act Article 50 Positioning

29.1 No categorical exemption claim

Article 50 applicability depends on whether the deployed functionality is a relevant AI system and on the operator's role as provider or deployer. A deterministic public-static simulation without a live generative model may present a narrower profile than a protected runtime using generative models, but the Office does not rely on a blanket conclusion that it is outside scope. Regulation (EU) 2024/1689 and current Commission guidance govern the analysis (EUROPEAN PARLIAMENT AND COUNCIL OF THE EUROPEAN UNION 2024; EUROPEAN COMMISSION 2026b).

29.2 Application date, transitional rule and enforcement posture

Article 50 applies from 2 August 2026. Commission materials describe a limited transitional period until 2 December 2026 only for the marking and detection duty applicable to relevant AI systems placed on the market before 2 August 2026 (EUROPEAN COMMISSION 2026c). The Office therefore does not base its posture on a categorical exemption. A protected deployment using a live interactive or generative model maintains a current provider/deployer analysis, disclosure before interaction, output provenance where applicable, human-review state and re-evaluation triggers.

29.3 Code-of-practice posture

The Commission and AI Board have treated the 2026 Code of Practice on Transparency of AI-Generated Content as capable of facilitating compliance, while noting that adherence is not conclusive proof of compliance (EUROPEAN COMMISSION 2026a). GoalOS may use such instruments as implementation evidence, not as a substitute for analysing actual conduct.

29.4 Compliance-by-design controls

The Office uses a conservative posture :

1. persistent disclosure that scenarios, scores and council deliberations are algorithmic outputs;
2. disclosure of runtime mode, model/provider where applicable and human-review status;
3. machine-readable provenance metadata in downloadable outputs when technically applicable;
4. no emotion recognition or biometric categorization features;
5. no deepfake-generation feature;
6. no automatic publication of public-interest text without accountable human review;
7. provider/deployer role memo for each protected deployment;
8. versioned Article 50 applicability record with owner and re-evaluation trigger.

TABLE 29.1 – Article 50 positioning matrix.

Office mode	Conservative position	Control
Public deterministic simulation	May or may not satisfy the legal definition depending on implementation; do not overclaim	Persistent algorithmic-output notice and no external execution
Interactive live-model runtime	Treat as potentially relevant interactive AI	Inform users before interaction; record model/runtime and human review
Synthetic text export	Treat marking/disclosure as potentially relevant	Provenance metadata, output label and export notice
Professional or public-interest publication	Human accountable review required	Named approver, evidence docket and publication receipt
Biometric / emotion / deepfake functions	Outside intended product scope	Feature prohibition and protected-deployment review

Office Law V — Transparency Survives Classification

Even where a particular Article 50 duty is ultimately determined not to apply, truthful algorithmic-output disclosure remains a product-integrity control unless a qualified reason justifies otherwise.

Québec Language Architecture

The Office provides complete, equal-quality French and English operating interfaces. French is available before access and contract acceptance; users can expressly choose English. Navigation, fields, statuses, dynamic messages, exports and legal terms are translated—not merely the Legal Center.

Québec commercial-language and adhesion-contract requirements inform this design (OFFICE QUÉBÉCOIS DE LA LANGUE FRANÇAISE 2026b; OFFICE QUÉBÉCOIS DE LA LANGUE FRANÇAISE 2026a). The release should maintain translation parity through a controlled string register and bilingual QA.

Privacy and Québec Law 25 Architecture

31.1 Data minimization

The public release should collect no confidential mission data by default and should explain browser-local storage, wallet interaction, hosting logs, RPC providers and user-controlled export. Protected deployments require a data map, purpose specification, retention/deletion schedule, access controls and contracts.

31.2 Privacy governance

The operator should publish the title and contact of the Privacy Officer, maintain governance policies, process access/correction requests, conduct privacy impact assessments where required and maintain an incident register (COMMISSION D'ACCÈS À L'INFORMATION DU QUÉBEC 2026c; COMMISSION D'ACCÈS À L'INFORMATION DU QUÉBEC 2026b; COMMISSION D'ACCÈS À L'INFORMATION DU QUÉBEC 2026a).

31.3 Automated decisions

The public Office does not make binding decisions about individuals. Any deployment involving decisions based exclusively on automated processing, or high-impact employment, credit, insurance, health, education or public-benefit contexts, is outside the public release and requires separate legal and impact review.

Contract and Publisher Architecture

The Legal Center should identify the registered operator, operating names, public contact, version, effective date, governing law, document bundle and release hash. Consent should bind the exact legal-bundle digest, language, origin, wallet or institutional principal, qualifying access right, issue time and expiry.

The terms should cover :

- research/experimental and no-advice boundary;
- conditional licence and prohibited uses;
- user responsibility for facts, authority and professional review;
- ownership, confidentiality and reuse;
- third-party, blockchain and provider risks;
- warranty disclaimer and non-reliance;
- liability allocation subject to mandatory law;
- no agency, fiduciary or employment relationship;
- suspension, termination, survival and severability;
- French-language delivery and express English choice;
- consumer and mandatory-law savings.

Protected-Deployment Triggers

The following conditions require a separately reviewed protected deployment :

- confidential, privileged, regulated, classified or personal information;
- persistent enterprise connectors or credentials;
- external transaction authority;
- money movement or custody;
- legal, tax, fiduciary, health or regulated professional decisions;
- decisions materially affecting an individual;
- critical infrastructure or physical-system control;
- high-consequence cybersecurity action;
- automated public-interest publication;
- deployment in jurisdictions requiring local representation, registration or impact assessment.

French Public Notice

Avis public canonique — français

Institution expérimentale de recherche, de simulation, de conception de missions, de gouvernance de la preuve et de surveillance institutionnelle. Les sorties sont algorithmiques et ne constituent ni un conseil juridique, fiscal, financier, médical ou professionnel, ni une autorisation, une approbation ou une garantie de résultat. Aucune exécution externe n'est effectuée par l'édition publique. Les faits réels, l'autorité responsable, les professionnels qualifiés et le droit impératif prévalent.

English Public Notice

Canonical public notice — English

Experimental research, simulation, mission-design, evidence-governance and institutional-monitoring system. Outputs are algorithmic and do not constitute legal, tax, financial, medical or other professional advice, authority, approval or a guarantee of results. The public edition performs no external execution. Actual facts, accountable authority, qualified professionals and mandatory law prevail.

Septième partie

Institutional Economics

The Office Value Function

The Office creates value by reducing missed opportunity, delay, coordination cost, proof cost, governance failure and institutional regret while improving accepted outcomes, resilience, option value and productive compounding.

Define verified Office value :

$$VOV = V_o + V_r + V_c + V_p + V_x - C_o - C_p - C_d - C_r - C_u,$$

where V_o is opportunity captured, V_r risk retired, V_c capability created, V_p productive capacity and V_x option value; costs include operation, professionals, delay, risk and unwind.

The Office does not book discovered opportunities as realized value. Every benefit progresses through explicit states :

identified → screened → validated → authorized
 → implemented → realized → independently verified.

Commercial Product Architecture

TABLE 37.1 – Illustrative Office product ladder; not a binding offer.

Edition	Illustrative annual price	Scope
Professional Office	US\$2,500–5,000	One operator ; browser-local missions and evidence
Team Office	US\$15,000–40,000	One legal organization ; shared operating discipline
Institution Office	US\$75,000–250,000	Governance, onboarding, source operations and executive cadence
Protected Office	US\$250,000–1,000,000+	Private runtime, connectors, security and mission execution
Portfolio Office	US\$1,000,000–5,000,000+	Multiple institutions or investment portfolio
Sovereign/strategic mandate	US\$5,000,000–25,000,000+	Public institution, infrastructure or regional programme

These values are management hypotheses. Pricing must reflect current scope, evidence, liability, professional authority and realized customer value.

AI-Native Operating Model

The Office should own intelligence, orchestration, proof, governance, customer relationship, Chronicle and outcome measurement. It should selectively route reserved legal, tax, regulated, audit, engineering and physical work to qualified providers.

This model avoids recreating a large professional firm on payroll while preserving accountable authority where AI cannot create it.

38.1 Cost discipline

AI can compress research, monitoring, drafting, comparison, simulation and reporting. It cannot make insurance, accountable professional authority, security, customer trust, physical execution and regulatory accountability free. Those costs should be explicit, mission-specific and generally customer-funded.

Recurring Revenue and Renewal

Renewal should be evidence-based. The Office must demonstrate :

- changes detected before the customer would likely have discovered them ;
- incomplete actions prevented ;
- opportunities validated and implemented ;
- Proof Debt retired ;
- architectures repaired or reconstituted ;
- realized value and risk reduction ;
- successor improvement ;
- continued relevance of the institutional twin.

Capital-to-Capacity Engine

Verified contribution may be allocated to :

- stronger models and evaluation;
- proprietary data and tools;
- security and resilience;
- human capability and professional authority;
- distribution and customer acquisition;
- compute, energy, laboratories and infrastructure;
- harder mission portfolios.

The larger compounding chain is :

intelligence → work → proof → authority → capability → value → capital → capacity.

Energy and compute constraints increasingly shape AI infrastructure decisions; the Office should evaluate physical capacity with complete capital, energy, permitting, resilience and externality denominators (INTERNATIONAL ENERGY AGENCY 2025).

The Productive-Capacity Frontier

The Office converts verified institutional value into durable capacity rather than allowing gains to dissipate into undifferentiated spending. Capacity may include data, distribution, security, expert authority, compute, laboratories, energy, grid access, land, facilities and industrial infrastructure.

The International Energy Agency estimates that data centres consumed approximately 415 TWh in 2024 and projects approximately 945 TWh by 2030 in its base case, with AI the most important driver of the increase (INTERNATIONAL ENERGY AGENCY 2025). The Office therefore treats compute as a coupled mission involving demand, site, grid, energy, capital, permits, equipment, resilience, community effects and verified productive output.

$$\text{PCV}(p) = V_{\text{output}} + V_{\text{option}} + V_{\text{strategic}} + V_{\text{public}} - C_{\text{capital}} - C_{\text{energy}} - C_{\text{delay}} - C_{\text{risk}} - C_{\text{externality}} - C_{\text{unwind}}.$$

A project is admitted only when demand, authority, evidence, financing, reversibility and public consequences meet the mission constitution.

Capacity doctrine

Verified intelligence should create the evidence, capital and institutional competence required to build greater productive capacity - while preserving the authority to stop, reconfigure or revoke what no longer deserves to scale.

Huitième partie

Empirical Authority and Production Proof

OfficeBench Ω

OfficeBench compares :

1. conventional multidisciplinary team;
2. static frontier-model report;
3. ungated agent workflow;
4. proof-governed Office mission;
5. Chronicle-informed successor.

Measures include :

- outcome quality and verified value;
- elapsed time and complete cost;
- human/reviewer burden;
- evidence completeness and currentness;
- authority compliance;
- failure containment and rollback;
- false recommendation rate;
- Chronicle precision;
- fresh-successor improvement;
- institutional regret.

Evaluation Integrity and Secure Measurement

The empirical programme assumes that models, agents and operators can adapt to visible tests. Evaluation therefore requires hidden cases, no-access baselines, context separation, evaluator rotation, pre-registered criteria, adversarial canaries and tamper-evident evidence. Where sensitive assets cannot be disclosed, privacy-preserving or secure evaluation may allow measurement without surrendering the underlying data, models or benchmarks (NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY 2026b).

Threat	Required control
Evaluation awareness	Hidden tasks, held-out criteria and delayed disclosure
Benchmark gaming	Fresh missions and independent outcome anchors
Evidence manipulation	Signed provenance, validator separation and raw-artifact access
Evaluator capture	Rotation, conflict disclosure and multi-party verdicts
Strategic underperformance	Canary missions and cross-model replication
Metric monoculture	Multi-objective scorecard and adverse-outcome register
Confidential benchmark leakage	Controlled enclaves or privacy-preserving evaluation

Office Law XIV — The Evaluator Must Also Be Governed

No independent-verification claim is stronger than the independence, authority, evidence access and conflict controls of the evaluator that issued it.

First Commercial Proof Run

The minimum external authority threshold is :

Gate	Minimum
Unrelated paying Office customers	3
Substantially comparable recurring missions	3
Independently verified customer outcome	1
Implemented route with measurable result	1
Separate repeat purchase or expansion	1
Active annual Office renewal	1
Rights-cleared Chronicle capability	1
Fresh-successor improvement	1
Material trust breaches	0
Unauthorized external actions	0

A public-safe Receipt Pack should state the paid scope, preregistered tests, evidence summary, independent verdict, complete cost denominator, implementation outcome, limitations, Chronicle decision and exact claim supported.

Monitoring Metrics

Metric	Governing question
Verified institutional value	What accepted value was actually realized and independently supportable?
Time-to-capability	How quickly did a frontier change become production-admissible capability?
Proof velocity	How quickly was decision-changing uncertainty retired?
Authority coverage	Are all consequential actions within current, attributable authority?
Mission autonomy	What completed without unplanned rescue, and at what risk?
Human burden	How much review, intervention and expertise was required?
Failure containment	Were errors detected, blocked and reversed?
Provider optionality	Can the institution switch models, tools and runtimes?
Chronicle precision	What admitted capability remains valid, and what was revoked?
Successor improvement	Did a fresh mission improve without governance regression?
Capital efficiency	How much verified capacity was created per dollar?
Large-scale impact	What effects occurred beyond the immediate institution?

1,000-Day Roadmap

46.1 Days 0–30 : constitute

Freeze the Office mandate, source constitution, authority model, Capability Twin, proof schemas and first mission portfolio.

46.2 Days 31–100 : prove

Run three bounded missions, one live architecture auction, one independent verdict and one Receipt Pack. Complete live access and incident exercises.

46.3 Days 101–365 : recur

Operate weekly/monthly/quarterly cadences, convert one customer to annual Office control, establish source refresh and prove one fresh successor.

46.4 Year 2 : productize

Add protected identity, connectors, source operations, provider auctions, role-based workspaces and benchmark publication.

46.5 Year 3 : scale

Expand Mission Packs, professional/validator network, industry constitutions and productive-capacity programmes while maintaining zero material unauthorized action.

Principal Risks and Controls

Risk	Failure mode	Control
Simulation theatre	Modeled success presented as realized outcome	State labels, external receipts and independent verification
Frontier hype	Vendor claim becomes institutional priority	Replication, mission relevance and production admissibility
False authority	Model output treated as professional or legal decision	Role classification, qualified authority, transaction gate
Proof overload	Costly evidence collected without decision value	Value-of-Proof queue
Metric gaming	Office optimizes dashboard indicators	Reality anchors, blind evaluation and adverse-outcome register
Memory poisoning	False or stale capability enters Chronicle	Admission predicate, expiry, revocation and fresh transfer
Founder dependence	Every exception returns to one person	Fixed products, delegated authorities, runbooks and decline paths
Provider lock-in	Institution cannot switch execution stack	Open schemas, model auction, portable state and fallbacks
Data leakage	Public edition receives sensitive facts	No-secrets gate, minimization and protected-deployment trigger
Regulatory overclaim	Disclaimer used as supposed exemption	Conduct-based role analysis and professional review
Infrastructure over-reach	Capital deployed before demand or permits	Milestone finance, SPVs, customer commitments and stop gates
Externality blindness	Local optimization creates public harm	Externalities auditor and large-scale impact review

Neuvième partie

Mission Casebook

Enterprise Operations Office

Objective : reduce operational friction while keeping service, security and authority measurable. The Office may coordinate onboarding, procurement, identity, finance, customer operations and incident response. The acceptance unit is not a closed ticket ; it is a verified completed mission with stakeholder acceptance and reusable evidence.

Illustrative mission : onboard a senior employee across payroll, identity, benefits, equipment, security training, expense and CRM systems by a fixed date. The Office creates typed dependencies, issues authority envelopes, verifies each system state and records acceptance.

Global Launch and Expansion Office

Objective : keep jurisdiction, entity, capital, talent, market, capability and infrastructure architecture current. The Office monitors funding openings, tax and treaty changes, permanent establishment, substance, ownership, hiring, capital, sanctions, procurement and infrastructure opportunities.

The result is not a static best-country ranking. It is a function-by-function architecture with evidence, professional authority, transaction gates and explicit reconstitution triggers.

Finance and Treasury Office

Objective : close, reconcile, collect, forecast and control with evidence. The Office begins with preparation, exception resolution, evidence and approvals—not unconstrained money movement or unsupervised filing.

A monthly close mission may retrieve ledgers, identify exceptions, request support, run reconciliations, model cash, draft explanations, route controller approvals and produce an audit-ready Evidence Docket.

Research Acceleration Office

Objective : move from hypothesis to independently supportable result while preserving reproducibility, safety and research rights. Jobs may include literature synthesis, experiment design, environment preparation, execution, replication, statistical review, artifact packaging and publication approval.

Chronicle admits only narrow methods whose provenance, rights and transfer conditions are explicit.

Public-Interest Mission Office

Objective : deliver public outcomes with legitimacy, transparency, accessibility and accountable authority. The Office must add public-purpose criteria, due-process and contestability controls, procurement rules, records management, accessibility, equality and large-scale impact review.

The public Office should not automate high-impact determinations about people without a deployment-specific legal and institutional constitution.

Professional Intelligence Office

Objective : turn research into a governed decision without impersonating licensed authority. The Office maps facts, retrieves current sources, identifies conflicts, generates scenarios, exposes Proof Debt, briefs qualified professionals, records opinions and blocks implementation until required receipts exist.

Portfolio Capital Office

Objective : compare mission opportunities across a portfolio using common evidence and verified-value states. The Office separates modeled, authorized, implemented, realized and independently verified value; only the latter can support claims of proven contribution.

Compute and Infrastructure Office

Objective : govern productive capacity across compute, energy, land, permits, network, equipment, financing, construction and commissioning. Every project requires customer commitments, milestone finance, technical authority, environmental and community impact review, stop conditions and verified productive output.

Dixième partie

Constitutional Appendices

Minimum Recurring Office Contract

A production contract should contain at least :

1. legal parties and accountable sponsor;
2. Office objective, scope and exclusions;
3. operating cadence and service levels;
4. data classes, sources, connectors and environments;
5. authority model and reserved decisions;
6. mission portfolio and value hypothesis;
7. proof obligations and independent evaluators;
8. professional-authority dependencies;
9. evidence, retention and deletion rules;
10. security, incident and continuity responsibilities;
11. rights, confidentiality, publication and reuse;
12. pricing, expenses and external-professional funding;
13. stop, pause, escalation, termination and rollback;
14. Chronicle and successor conditions;
15. governing law, mandatory-law savings and language mechanics.

Authority Envelope Schema

```
{
  "envelope_id": "ae-...",
  "principal": {"type": "institution", "id": "..."},
  "actor": {"type": "agent|human|provider", "id": "..."},
  "mission_id": "mission-...",
  "purpose": "bounded description",
  "allowed_tools": [],
  "allowed_data_classes": [],
  "allowed_actions": [],
  "prohibited_actions": [],
  "budget": {"currency": "USD", "max": 0},
  "effective_at": "ISO-8601",
  "expires_at": "ISO-8601",
  "required_evidence": [],
  "acceptance_tests": [],
  "escalation": [],
  "rollback": {},
  "revocation": {},
  "approvals": []
}
```

Evidence Docket Schema

Object	Required content
Objective record	Sponsor, decision, value at stake, deadline, success criteria
Fact register	Fact, official source, retrieval date, owner, confidence, expiry
Assumption ledger	Assumption, sensitivity, evidence needed, decision impact
Claims matrix	Material claim, evidence, contradiction, status and blocker
Job ledger	Actor, tools, data, budget, output, failure and verifier
Cost ledger	Models, data, professionals, support, delay and opportunity cost
Authority register	Approval, scope, owner, effective date, expiry and revocation
Action receipts	Transactions, filings, payments, deployments and acceptance
Outcome record	What happened, value, adverse effects and delayed consequences
Rights register	Customer, platform and third-party rights; reuse conditions
Chronicle decision	Admit, limit, repair, reject, revoke, expire or supersede
Successor result	Baseline, fresh mission, metrics, regressions and verdict

Chronicle Capability Record

```
{
  "capability_id": "cap-...",
  "statement": "narrow capability",
  "scope": [],
  "prohibited_generalizations": [],
  "originating_missions": [],
  "evidence_refs": [],
  "acceptance_refs": [],
  "rights_basis": [],
  "known_failures": [],
  "model_tool_versions": [],
  "effective_at": "ISO-8601",
  "expires_at": "ISO-8601",
  "revocation_conditions": [],
  "fresh_transfer_tests": [],
  "status": "candidate|admitted|limited|expired|revoked|superseded"
}
```

Source Pack Schema

A signed source pack contains source URL, publisher, title, jurisdiction, source class, effective date, retrieval time, content hash, language, authority level, affected claims, expiry/recheck condition and publisher signature. A source pack establishes provenance and currentness of the captured source, not infallibility of interpretation.

Incident Register

Minimum fields include incident identifier, discovery time, reporter, affected systems and people, data classes, mission and authority state, containment, evidence, severity, notification analysis, responsible owner, corrective action, Chronicle impact and closure approval.

Weekly Office Brief Template

1. What materially changed?
2. Which assumptions, evidence or authority are affected?
3. Which missions are blocked or newly admissible?
4. What opportunity has the highest Value of Proof?
5. Which decision is due before the next brief?
6. What should stop, pause, continue or accelerate?
7. Which Chronicle records require revalidation?
8. What verified value was realized?

Monthly Architecture Auction Agenda

1. confirm mission constitution and material events ;
2. refresh Capability Twin and source state ;
3. regenerate candidate architectures ;
4. run counterfactual stress worlds ;
5. review Proof Debt and Value-of-Proof queue ;
6. select immediate, successor and reserve architectures ;
7. approve bounded jobs and authority envelopes ;
8. review value, resilience, externalities and rollback ;
9. record acceptance and board decisions.

Article 50 Applicability Record

The deployment file should document :

- feature and user interaction description;
- whether a live AI model is used;
- provider, deployer, importer or distributor role analysis;
- whether outputs include synthetic text, audio, image or video;
- whether content is public-interest information;
- whether biometric categorization or emotion recognition exists;
- user disclosure, machine-readable marking and human-review controls;
- owner, counsel, date, evidence and re-evaluation triggers.

Legal Launch Checklist

Control	Owner	Status
Registered operator and contact identified	Publisher	CONDITION
Complete French and English operating editions	Product	CONDITION
Exact legal bundle and consent evidence	Legal/Product	CONDITION
Privacy Officer and data map	Privacy	CONDITION
Incident register and response	Security/Privacy	CONDITION
Article 50 role/applicability record	Legal / AI Governance	CONDITION
Public-claims evidence register	Communications	CONDITION
Live access acceptance test	Operations	CONDITION
Protected-deployment contracts	Legal/Security	CONDITION
Professional-authority network	Operations	CONDITION
Insurance bound for actual conduct	Executive	CONDITION
Independent deployment-specific legal review	Counsel	CONDITION

Glossary

Accepted outcome. A result that has passed its evidence and accountable acceptance gate.

Authority Envelope. The scoped, minimal and revocable permission under which an actor may work.

Chronicle. The governed memory layer that admits, versions, expires and revokes reusable capability.

Evidence Docket. The complete, attributable record supporting a consequential decision or outcome.

Institutional Pulse. The board-readable summary of Office state and next decision.

Proof Debt. Unresolved claims, stale evidence, missing authority or untested conditions capable of changing a decision.

Recurring control plane. The continuous system that monitors, governs, proves and reconstitutes an institution.

Singularity Exposure Ratio. Institutional adaptation lag divided by the half-life of mission-relevant assumptions.

Verified value. Value progressed through evidence-backed states to actual realization and, where claimed, independent verification.

Value of Proof. The expected decision impact of obtaining a specific evidence object relative to its cost, delay and risk.

AI-Generated Visual Asset Record

This edition includes selected image-generation visual-direction plates. The controlled record identifies the asset, generation date and editorial purpose. Generated images are not evidence of external facts and do not override exact LaTeX text, equations, tables, schemas or cited sources.



Author and Institutional Stewardship

Vincent Boucher is President of MONTREAL.AI and QUEBEC.AI. This publication is issued by MONTREAL.AI Research as a founding institutional design, operating and commercialization paper.
Public contact : info@quebec.ai.

Recommended citation

Boucher, Vincent. 2026. *GoalOS Singularity Office Ω : The Recurring Control Plane for a Living Institution*. Definitive Production Institution Edition, v3.0.0-SO3. MONTREAL.AI Research.

ImageGen Institutional Visual Atlas

Visual-status notice

The following plates were generated as institutional visual-direction studies and integrated under editorial control. They are illustrative rather than normative. Exact LaTeX text, equations, tables, schemas and cited sources control wherever a plate differs from the authoritative publication.



GoalOS Singularity Office Ω

The Recurring Control Plane for a Living Institution

Tools to Ride the Singularity

Vincent Boucher

President, MONTREAL.AI & QUEBEC.AI

Montréal, Québec, Canada | 27 July 2026

EXECUTIVE SUMMARY

*the Navigator determines
where to go;
the Office ensures the
institution keeps deserving
to go there.*



Abstract

This paper introduces the GoalOS Singularity Office (Ω) as the recurring control plane for a living institution. The Office operates a complete institutional loop—Observe, Govern, Prove, Reconstitute, Compound—that preserves merit, compounds capability, and aligns every action with sovereign purpose. It transforms strategy into enforceable, self-reinforcing operation, so the institution not only moves in the right direction but remains worthy of the future it seeks to create.

Why an Office is needed

Ambition without control decays. Strategies drift, incentives corrode, and capabilities stagnate when no system ensures continuous integrity. The Office exists to close that gap by making institutional quality a recurring outcome rather than an episodic initiative. It encodes sovereign discipline into operational practice and installs a memory of merit that outlives individuals and cycles.

Core thesis

A living institution endures by recurring control, not by static plans. The Office is the minimal, sufficient control plane that observes reality, governs choices, proves results, reconstitutes structure, and compounds capital—relentlessly and measurably.

Contributions

We define the Office Principle and its five-step loop, specify the control plane architecture, introduce merit-preserving governance patterns, and provide reference metrics and interfaces for implementation. Together, these contributions form a portable blueprint for sovereign institutions operating at civilizational scale.

Paper structure

Section 1 motivates the control-plane requirement. Section 2 defines the Office and the five-step loop. Section 3 details each step and its instruments. Section 4 presents the architecture and data model. Section 5 reviews metrics and fiduciary guardrails. Section 6 illustrates deployment and integration. Section 7 concludes with future work.

Office Principle

The Office operates a closed institutional loop that converts observation into compounded sovereign capability.

- 1 OBSERVE**
Sense reality continuously across all domains that matter.
- 2 GOVERN**
Decide with sovereignty; aligning resources to purpose.
- 3 PROVE**
Measure truth, verify outcomes, and establish merit.
- 4 RECONSTITUTE**
Adapt structures, roles, and rules to preserve integrity and fit.
- 5 COMPOUND**
Reinvest capability and knowledge to increase institutional power.

Repeat endlessly. Improve relentlessly.

THE OFFICE LOOP



A CONTINUOUS LOOP.
A LIVING INSTITUTION.

OBSERVE + GOVERN + PROVE + RECONSTITUTE + COMPOUND

2. Architecture of GoalOS Singularity Office Ω

A Constitutional Control Plane for Perennial, Living Institutions

GoalOS Singularity Office Ω is the institutional control plane that converts discovery into recurring impact. It encodes mission, authority, memory, and recurrence into a closed loop of intelligence and governance—designed to endure, compound, and remain deserving of delegated power across generations.

2.1 Constitutional Architecture

The Office is grounded in a written constitution and operating standards that define purpose, powers, boundaries, and fiduciary duties. All actions must be consistent with the mission, lawful, and aligned with long-term institutional value.

- **Mission First:** every decision advances the mission.
- **Bounded Authority:** powers are explicit, minimal, and revocable.
- **Accountability:** actions are observable, provable, and reviewable.
- **Perpetuity:** institutions outlive individuals.

2.2 Operating Loop

The Office runs a perpetual control loop that senses, reconstitutes, executes, proves, and compounds. Each cycle strengthens institutional deservingness and expands the frontier.

- Sense the frontier and detect material changes.
- Reconstitute mission and priorities.
- Exercise bounded authority with required approvals.
- Execute through AGI jobs and toolchains.
- Produce proofs, obtain validators, and record.
- Accept, release, and compound the mission.

2.3 Evidence and Memory

Evidence is the currency of institutional legitimacy. All material actions are accompanied by structured proofs and stored in the Chronicle—an immutable, queryable memory that enables learning, audit, and continuity.

- **Proof Bundles:** artifacts, metrics, and attestations.
- **Validators:** independent or internal attestation.
- **Chronicle Memory:** tamper-evident, time-indexed.
- **Retrieval & Learning:** informs better reconstitution.

GOALOS SINGULARITY OFFICE Ω — ARCHITECTURE



2.4 Chronicle and Successor Rights

The Chronicle preserves institutional memory and establishes successor rights so that future stewards inherit not only assets, but context, obligations, and proven ways of creating value.

- **Continuity of Mission:** purpose transcends persons.
- **Successor Rights:** clearly defined inheritance of authority and duties.
- **Compounding Value:** each generation starts further ahead.
- **Institutional Immortality:** the Office endures.

NAVIGATOR vs. OFFICE

NAVIGATOR <i>Defines Direction</i>		OFFICE Ω <i>Maintains Institutional Deservingness</i>
Sets long-term vision and strategic intent.		Encodes mission into rules, standards, and operating systems.
Explores opportunities and unknowns.		Governs execution and recurrence within bounds.
Takes calculated risks to open new frontiers.		Requires evidence and approval to act.
Asks, "What could be?"		Asks, "What is right, true, and provably beneficial?"
Discovers and proposes new possibilities.		Delivers and compounds real, measurable value.

MATHEMATICAL NOTATION

	Recurrence $S_{t+1} = \mathcal{R}(S_t, E_t, A_t, C_t, \Theta_t)$ State at time $t + 1$ is a function of state, evidence, actions, capital, and constraints.
	Evidence $E_t = \{e_i, v_i, \mu_i, \sigma_i\}_{i=1}^n$ Evidence set E_t contains items e_i with validators v_i , metrics μ_i , and confidence σ_i .
	Compounding $V_{t+1} = V_t + \Delta V_{\text{real}} - \Gamma_t + \Phi(E_t)$ Value grows by realized value, minus friction, plus value uplift from evidence-derived learning.

4

Governance, Public-Safe Deployment, Legal Perimeter, and Institutional Economics



Governance and bounded authority

Singularity Office Ω is governed by a layered model of bounded authority. The system executes the governance cycle only within the scope explicitly delegated by the institution. Authority is compartmentalized as Observe → Govern → Prove → Reconstitute → Compound, with each transition producing verifiable evidence. No autonomous control is permitted outside the declared perimeter. Human oversight remains ultimate.



Public-safe operating perimeter

The public edition is designed for browser-local or public-static deployment. It requires no installation, no persistent identity, and no user secrets. All activation is user-controlled and revocable. The system processes inputs in-session only, persists public artifacts to immutable logs, and exposes no hidden channels. Evidence logs are append-only and exportable. Operators retain full custody of their environments.



Regulatory posture and legal positioning

GoalOS Singularity Office Ω is positioned as governance support infrastructure, not as a provider of professional, regulated, or fiduciary services. The public edition does not automate professional judgments, regulated determinations, or fiduciary decisions. Outcomes are advisory, explainable, and always subject to human review and institutional policy. Operator responsibility is explicit, separated, and non-delegable.



Institutional economics

Singularity Office Ω converts governance into measurable, recurring value. The control plane reduces institutional drift, accelerates adaptation, hardens proof discipline, and compounds reusable capability. Costs shift from bespoke projects to governed recurrence, enabling transparent stewardship and defensible returns over time.

COMPARATIVE POSITIONING

	LEGACY AUTOMATION	GENERIC COPILOTS	GOALOS SINGULARITY OFFICE Ω
AUTHORITY	Centralized, rigid workflows	Prompt-driven, user-scoped	Bounded, institutional, and auditable
EVIDENCE	Fragmented, siloed artifacts	Ephemeral, hard to verify	Structured, timestamped, verifiable
ACCEPTANCE	Low (black-box outcomes)	Variable, trust by proximity	High (explainable, policy-aligned)
MEMORY	Application-specific silos	Context-limited, non-persistent	Institutional memory & knowledge graph
RECURRENCE	One-off projects, rework	Ad-hoc, inconsistent reuse	Recurring control plane & compounding

RECURRING VALUE LEVERS

Lower institutional drift	Continuous alignment of policy, operations, and evidence.
Faster adaptation	Earlier detection, scenario rehearsal, and controlled response.
Better proof discipline	Evidence-native operations that withstand scrutiny.
Reusable capability	Modular playbooks and patterns across functions and cycles.
Governance visibility	End-to-end transparency for boards, regulators, and stakeholders.
Compounding institutional memory	Knowledge accrual that strengthens decisions over time.



PUBLIC EDITION PRINCIPLES (NON-NEGOTIABLE)

Browser-local or public-static deployment

No professional, regulated, or fiduciary automation

No user secrets required for the public edition

User-controlled activation and revocable use

Conditional licensing for access and features

Evidence logging with export and auditability

Separated operator responsibility

5. Use Cases and Mission Classes

The Singularity Office applies a common control logic across diverse institutional contexts. Each mission class maps recurring institutional frictions to operational patterns the Office is designed to observe, govern, prove, reconstitute, and compound.

MISSION CLASS	PRIMARY FRICTION	CONTROL FOCUS	TYPICAL TRIGGERS	PRIMARY OUTCOMES
 Governance Refresh	Fragmented accountability	Charter clarity, rights, duties, and cadence	Mandate change, leadership transition	Decision velocity, trust, legitimacy
 Compliance Posture	Regulatory drift and complexity	Policy coherence, controls, evidence flows	New regulation, audit findings	Lower friction, clean audits, regulatory trust
 Talent and Staffing	Skill gaps and misalignment	Role design, flows, incentives, enablement	Scaling, turnover, capability gaps	Retention, velocity, capability depth
 Treasury and Capital Discipline	Capital leakage and opacity	Allocation logic, guardrails, evidence	Budget reset, market stress	Capital efficiency, liquidity resilience
 Infrastructure Oversight	Siloed systems and fragility	Reliability, observability, resilience	Outages, obsolescence, scaling thresholds	Continuity, performance, recovery speed
 Market-Entry Adaptation	Local context and speed	Regime mapping, entry patterns, partnerships	New jurisdiction, expansion mandate	Lower friction entry, local legitimacy, revenue ramp
 Institutional Repair	Crisis, distrust, erosion	Stabilize, restore, prove, and rebuild legitimacy	Crisis event, erosion of trust	Stability, restored legitimacy, durable capacity

Mission classes are not mutually exclusive. The Office sequences and composes them per institutional priority.

7. Evaluation Programme

Every mission is evaluated on an evidence-first framework. Validators review evidence packs and score outcomes across six criteria to determine advancement, reuse, and compound potential.

CRITERION	QUESTION	MEASURE (EXAMPLES)	SCORE (1-5)
 Mission Turnaround	Did the mission achieve the intended outcome?	Time-to-outcome, KPI delta, benefit realization	★★★★★
 Evidence Completeness	Is the evidence sufficient, verifiable, and reproducible?	Artifacts present, lineage clarity, replicability	★★★★★
 Rollback Discipline	Could the mission be reversed safely?	Rollback plan quality, rehearsal, measured reversibility	★★★★★
 Validator Agreement	Do independent validators concur with results?	Validator set size, agreement level, confidence notes	★★★★★
 Accepted Reuse	Has the pattern been adopted elsewhere?	Number of adopters, contexts, adaptation notes	★★★★★
 Successor Improvement	Does the next iteration outperform this one?	Delta in outcomes, efficiency, resilience, or trust	★★★★★

Missions advancing to the next phase average ≥ 4.0 with no criterion < 3.0 .

6. Implementation Roadmap

The Office grows its mandate through constrained experimentation, evidence accumulation, and governance ratification. Each phase is reversible until proven.



 Each phase advances only when evidence thresholds are met and rollback remains clean.

8. Conclusion

Institutions do not fail for lack of plans; they fail for lack of recurring control. Static plans calcify. Context shifts. Incentives drift. Evidence thins. The institution slows—then breaks.

The GoalOS Singularity Office Ω provides the living control plane institutions need: observe precisely, govern cleanly, prove continuously, reconstitute safely, and compound relentlessly. It converts volatility into learned capability and uncertainty into institutional memory. It shortens the path from mandate to outcome and from outcome to enduring capacity.

This is how institutions remain sovereign in a world that moves faster than legacy cycles can follow.

The GoalOS Singularity Office Ω is not merely an assistant layer; it is a recurring constitutional operating layer for institutions that must adapt faster than static planning cycles can keep up. It is the institutional control plane for a living future.



Vincent Boucher
President, MONTREAL.AI & QUEBEC.AI
Montréal, Québec, Canada • 27 July 2026

N.1 Visual provenance

The visual system uses midnight navy, institutional gold, ivory, planetary networks, orbital control geometry and the Ω seal to represent a recurring constitutional institution rather than a conventional application dashboard. Exact typography and formal institutional architecture remain vector-typeset in LaTeX.

References

- COMMISSION D'ACCÈS À L'INFORMATION DU QUÉBEC (2026a). *Collecte de renseignements personnels par une entreprise*. URL : https://www.cai.gouv.qc.ca/protection-renseignements-personnels/information-entreprises-privées/collecte-renseignements-personnels_entreprises (visité le 27/07/2026).
- (2026b). *Incidents de confidentialité et mesures de sécurité pour les entreprises*. URL : <https://www.cai.gouv.qc.ca/protection-renseignements-personnels/information-entreprises-privées/incidents-confidentialite-mesures-securite-entreprises> (visité le 27/07/2026).
- (2026c). *Responsable de la protection des renseignements personnels dans une entreprise*. URL : <https://www.cai.gouv.qc.ca/protection-renseignements-personnels/information-entreprises-privées/responsable-protection-renseignements-personnels-entreprise> (visité le 27/07/2026).
- ETHEREUM NAME SERVICE (2026). *ENS Contract Deployments*. URL : <https://docs.ens.domains/learn/deployments/> (visité le 27/07/2026).
- EUROPEAN COMMISSION (9 juill. 2026a). *Commission Opinion on the Assessment of the Code of Practice on Transparency of AI-Generated Content*. URL : <https://digital-strategy.ec.europa.eu/en/library/commission-opinion-assessment-code-practice-transparency-ai-generated-content> (visité le 27/07/2026).
- (2026b). *Guidelines on Transparency Obligations for Certain AI Systems under Article 50 of the AI Act*. URL : <https://digital-strategy.ec.europa.eu/en/policies/guidelines-transparency-ai-generated-content> (visité le 27/07/2026).
- (2026c). *Transparency Obligations under Article 50 of the AI Act : Frequently Asked Questions*. URL : <https://digital-strategy.ec.europa.eu/en/faqs/transparency-obligations-under-article-50-ai-act> (visité le 27/07/2026).
- EUROPEAN PARLIAMENT AND COUNCIL OF THE EUROPEAN UNION (2024). *Regulation (EU) 2024/1689 Laying Down Harmonised Rules on Artificial Intelligence*. URL : <https://eur-lex.europa.eu/eli/reg/2024/1689/oj> (visité le 27/07/2026).
- INTERNATIONAL ENERGY AGENCY (2025). *Energy and AI*. URL : <https://www.iea.org/reports/energy-and-ai> (visité le 27/07/2026).
- NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. Rapp. tech. NIST AI 100-1. U.S. Department of Commerce. URL : <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf> (visité le 27/07/2026).
- (2026a). *AI Agent Standards Initiative : Interoperable and Secure AI Agents*. URL : <https://www.nist.gov/artificial-intelligence/ai-agent-standards-initiative> (visité le 27/07/2026).
- (27 mars 2026b). *CAISI and OpenMined Collaboration to Enable Secure AI Evaluations*. URL : <https://www.nist.gov/news-events/news/2026/03/announcement-caisi-signs-crada-openmined-enable-secure-ai-evaluations> (visité le 27/07/2026).
- (2026c). *Challenges in Monitoring Deployed AI Systems*. URL : <https://www.nist.gov/news-events/news/2026/03/new-report-challenges-monitoring-deployed-ai-systems> (visité le 27/07/2026).
- (2026d). *Identity and Authority for Software Agents*. URL : <https://www.nist.gov/news-events/news/2026/02/new-concept-paper-identity-and-authority-software-agents> (visité le 27/07/2026).
- OFFICE QUÉBÉCOIS DE LA LANGUE FRANÇAISE (2026a). *Contrats d'adhésion*. URL : <https://www.oqlf.gouv.qc.ca/francisation/entreprises/contrats-adhesion.html> (visité le 27/07/2026).
- (2026b). *Langue du commerce et des affaires : sites Web et publications commerciales*. URL : https://www.oqlf.gouv.qc.ca/francisation/droits_linguistiques/droits/langue-du-commerce-et-des-affaires.html (visité le 27/07/2026).
- OPENAI (2026). *The Next Evolution of the Agents SDK*. URL : <https://openai.com/index/the-next-evolution-of-the-agents-sdk/> (visité le 27/07/2026).
- OWASP FOUNDATION (2025). *OWASP Top 10 for Large Language Model Applications*. URL : <https://owasp.org/www-project-top-10-for-large-language-model-applications/> (visité le 27/07/2026).
- RAO, Anita, Andrew KELLER, Neha KALRA, Ryan STEED, Kweku KWEGYIR-AGGREY, Kevin KLYMAN, Diane STAHELI et Amanda BERGMAN (2026). *Challenges to the Monitoring of Deployed AI Systems*. Rapp. tech. NIST AI 800-4. National Institute of Standards et Technology. DOI : 10.6028/NIST.AI.800-4. URL : <https://doi.org/10.6028/NIST.AI.800-4> (visité le 27/07/2026).

WORLD WIDE WEB CONSORTIUM (2013). *PROV-O : The PROV Ontology*. URL : <https://www.w3.org/TR/prov-o/> (visité le 27/07/2026).

Canonical Operating Doctrine

Intelligence proposes. Authority bounds. Work executes. Evidence proves. Acceptance decides. Chronicle governs what survives. Capital expands what proved valuable.

**The Navigator determines where to go.
The Office ensures the institution keeps deserving to go there.**