

MONTREAL.AI · QUEBEC.AI

SUCCESSOR Ω USER GUIDE Ω

FROM ONE MISSION TO ENDURING INTELLIGENCE

IDEAS
EVIDENCE
PEOPLE
INSTITUTIONS
WORLDS
A BRIGHTER
HUMAN FUTURE

Ω
SCIENCE
GAME THEORY
ENGINEERING
ECONOMICS
SOVEREIGNTY
HUMANITY

Ω
LEARN
PROVE
GOVERN
COMPOUND
CIVILIZE

A MORE
INTELLIGENT
CIVILIZATION

Vincent Boucher

President, QUEBEC.AI & MONTREAL.AI

ILLUSTRATED AUTHOR EDITION 4.1 | PRODUCT 20.0.1

Conceptual artwork. Local A0 research. No production authority.

SUCCESSOR Ω USER GUIDE Ω

From one mission to enduring intelligence.

Vincent Boucher

President, QUEBEC.AI & MONTREAL.AI

Reader's compact

Begin with an observable decision. The first useful result is a precise mission, a candidate trained from permitted evidence, an honest comparison and an exact record you know how to restore.

Three workbenches, one discipline. Classification, predictive worlds and shared-resource compositions remain local research tools. The guide consistently distinguishes product functions, teaching tools and conditional analysis.

Six artworks, no invented measurements. The plates express an imagined horizon. Their decorative inscriptions do not describe an available interface, a result, proof or permission. Commands and technical figures control the usage described.

Two versions to distinguish. Guide 4.1; unchanged software 20.0.1 R10.1. This edition does not change contracts or qualification scope. Historical trials and newly executed checks have separate records.

Choose the mission. Earn the comparison.

Preserve what survives.

Contents

1	Start here: turn one mission into a retained intelligence asset	
2	What you own, what you run, and what stays human	
3	Choose the right workbench	
4	Install once, and keep the product distinct from your work	
5	Your first session: a complete synthetic cycle	
6	Configure a real mission before supplying real data	
7	Before formation: use SEIZE to choose the next decision	
8	Constitute the Mission Gym without confusing it with a world model	
9	Bind intention to the exact executable plan	
10	Turn the mission into twenty-one bounded work-and-proof jobs	
11	Preserve the distinct protocols, symbols and decisions	
12	Workbench A: customer-specific classification	
13	Read comparisons without fooling yourself	
14	Workbench B: executable predictive worlds	
15	Interpret uncertainty, contradictions and revision	
16	Workbench C: compose intelligence under shared constraints	
17	Pin, inspect and propose	
18	Compare compositions and decide locally	
19	Create a successor deliberately	
20	Restore an exact record without inventing permission	26
21	Troubleshooting without weakening the boundary	27
22	Academy, support and frequently confused features	28
23	A practical operating cadence	29
24	From a local specialist to an admitted institution	29
25	Use SUCCESSOR Ω for the highest useful value	30
26	A portfolio that can grow without outrunning evidence	31
27	Advanced operator note: prediction, optimization and value	34
28	Advanced operator note: scarcity and strategic behavior	34
29	Analytical laboratory: test the decision logic itself	35
30	Advanced operator note: physical accounts before large-scale promises	38
31	Physical integrity: a small exact system before an enormous horizon	39
32	Advanced operator note: productive expansion needs a network	41
33	Productive expansion: the threshold is conditional, delivery is real	42
34	Command reference: copy a complete command, then inspect its result	44
35	Evidence, provenance and what was tested for this guide	45
36	Glossary	47
References		49

PART / PARTIE I

First mission

1 Start here: turn one mission into a retained intelligence asset

The first useful outcome is not a grand label. It is one decision problem made explicit, one candidate trained from permitted evidence, one honest comparison, and one exact record you can reopen. This handbook shows how to obtain that outcome with SUCCESSOR Ω 20.0.1 R10.1.

Your durable work can include the mission, evidence mapping, competing hypotheses, trained parameters, executable programs, comparison records and reasons for retaining or rejecting a successor. Those objects can matter beyond a single session. They are not, by themselves, independent proof, a production controller or ownership of the underlying platform. [1, 2, pp. 10–12, 36–37]

1.1 Choose a reading path

You are here	Begin with	Finish with
First-time operator	Installation, then your first session	An inspectable synthetic run and restoration
Owner choosing a mission	SEIZE, mission dossier, value	A bounded evidence budget and explicit stop rule
Technical builder	Workbench selection, then the relevant route	A frozen candidate, limitations and eligible comparison
Reviewer or custodian	Claim interpretation, continuity, evidence register	A traceable decision, not a collection of green badges
Advanced reader	Control, games, physical accounts, new analytical laboratory	Explicit assumptions and reproducible conditional calculations

Best first move: run the optional teaching companion in this guide, or open the product’s integrated mission wizard. Use public synthetic evidence first. Neither route requires a wallet, model subscription or connection to production equipment. Technical installation is still required.

1.2 What this illustrated edition changes

This is **Guide 4.1 — Illustrated Author Edition**, accompanying the unchanged 20.0.1 R10.1 product. It retains the complete Guide 4.0 operating manuscript, command contracts, mathematical companion and teaching software. The new composition places six newly supplied conceptual artworks around the learning journey; none of the five style-reference pictures is reused.

The 4.0 execution campaign remains separately identified as historical evidence. The 4.1 records report only checks actually performed for this edition. This publication is not a product upgrade, a fresh industrial trial or a new grant of rights.

Read the image plates as an editorial horizon. Their words and diagrams are not measurements, working interfaces, approvals or implemented functions. Exact instructions, numbered technical figures, equations and source records govern the substance.

2 What you own, what you run, and what stays human

Think of the product as a **mission laboratory with a durable record**. You choose the objective, evidence, comparisons and restrictions. The software performs bounded learning and computation. You inspect the resulting candidates and decide what to retain.

A *mission* states one decision problem. *Evidence* is the data you are permitted to use, with its origin and exposure history. A *candidate* is a fitted model or executable hypothesis. A *composition* adds the planner, resource contract, objective, calibration and scope around a predictor. A *Chronicle* records ordered events. A *proof request* describes a trial that should be commissioned; it is not the resulting independent proof.

The customer retains mission data, configurations and generated mission objects within the contractual rights boundary. The platform is licensed rather than transferred. An export does not override third-party rights, make a self-declared source authentic, or grant ownership of an external model. [3, 2, pp. 36–37]

2.1 Four reading labels

Guide companion identifies optional teaching tools outside the product. They prepare synthetic lessons and call the original CLI; they add no production authority or new learning capability. **Product behavior** identifies functions in the inspected 20.0.1 archive. **Recommended practice** identifies a proposed operating discipline or worksheet in this guide. **Conditional theory** identifies mathematical models that clarify choices but are not new 20.0.1 features. These categories appear throughout the advanced chapters.

2.2 A0 is the boundary of every route

The words *recommend*, *reversible_experiment*, *produce* and *service* can occur as model labels or hypothetical proposals. They are not approved orders or permissions. 20.0.1 does not connect this guide to machinery, accounts, purchasing or production tools. A company using a proposal outside the software would make a separate, consequential decision for which this guide supplies no authorization.

The architectural design reserves *Specialist ASI* for mission-specific superiority with the required hard gates, proof and admission. A `.successor.json` suffix, a larger generation number or a good training score does not confer that status. [2, pp. 11–12, 33, 51–54]

CHOOSE THE REPRESENTATION

Classify 6 normalized signals 3 proposed classes Explicit snapshots	World Regular episodes Interpreted program Bounded predictions	Compose Predictor + planner 1–6 process cells Resources and scope
---	--	---

All routes: local A0 research, no automatic admission.

Figure 1. Choose by evidence representation, not by ambition. These are product routes, not proof levels.

3 Choose the right workbench

Start with the evidence you actually have—not the most impressive menu title.

Route	Required representation	What 20.0.1 learns or computes	What it does not establish
Classifier	30–1,000 cases, six normalized observations, preferred label and positive weight	Competing small neural classifiers, local comparison and deliberate succession	Arbitrary document understanding, causal action value or general business competence
World Engine	Regular, complete temporal episodes with current observation, executed action and later outcome	Typed transition programs, fitted parameters, structural edits and bounded planning	A universal digital twin, arbitrary hidden-state discovery or real causal identification
Composition	A compatible process predictor or temporal evidence, plus a one-to-six-cell contract	Neural fitting or model import, shared-resource allocation, identity and applicability checks	Plant control, general job-shop scheduling, a strategic auction or automatic financial Alpha
Academy	No private business evidence needed	Reading pathway, saved progress and source materials	Capability certification or an earned designation

The classifier’s six signals are evidence strength, outcome gap, intervention fit, stakeholder risk, reversibility and drift risk. They are not six arbitrary sensor columns. The composition adapter specifically represents normalized stored heat/fouling and ambient load/feed quality; it is not made suitable for a different industry by changing a title. [3]

When no route fits: write the intended states, actions, timing, outcomes and resource constraints in the mission worksheet. Narrow the mission or commission a separately evaluated adapter. Do not force incompatible evidence into a schema merely to make a run complete.

4 Install once, and keep the product distinct from your work

4.1 Three folders, three purposes

The complete guide archive contains 06_PRODUCT/ with the exact customer ZIP. Extract that ZIP into a directory you control. The extracted **product root** is the folder containing VERSION.json, VERIFY_PACKAGE.py and 07_NATIVE/.

work/	
USER_GUIDE_4_1/	handbook, source and teaching companion
SUCCESSOR_OMEGA/	unchanged extracted customer product
first-mission/	NEW practice folder, created by the companion

Do not work inside an unopened ZIP. Do not choose the seller console as the product root. Do not put experimental outputs inside the product: its integrity inventory is intended to remain stable. Retain the original ZIP before changing anything around it.

The optional 08_TOOLS/unpack_product.py extracts the included customer ZIP after checking its exact identity and restores stored executable modes. It requires a **new** destination. It is an extraction helper, not a software installer, runtime downloader or elevation tool.

4.2 Check your runtime without mistaking it for qualification

The native console runs on Node.js; Python is used for package verification and guide utilities. Check the versions available on your computer:

```
node --version
python3 --version
```

The source Guide 4.0 campaign was executed on Linux x86_64 with **Node.js 22.16.0** and **Python 3.13.5**. That is a reproducibility fact, not a recommendation to install an older patch. Official Node documentation recommends Active or Maintenance LTS for production applications; the official download page was checked for Guide 4.0. Current-patched-LTS, Windows and macOS operation were **not** qualified here. A current security patch and a successful first-session compatibility check are separate requirements. [4]

Where python3 is not the command installed on your computer, resolve the correct interpreter with your administrator. The one-line command catalogue avoids POSIX continuation characters, but it is not a claim of Windows qualification. Do not disable administrator policies to reproduce this guide.

4.3 Verify the actual product

Open a terminal in SUCCESSOR_OMEGA and run:

```
python3 VERIFY_PACKAGE.py
node 07_NATIVE/start.mjs doctor
node 07_NATIVE/start.mjs mission wizard
```

A successful inventory means that the extracted contents match the distribution's recorded identities and modes. It does not establish the truth of customer evidence or suitability for a real mission. doctor reports the runtime and workbench boundary; it is not a security audit.

Choose en or fr. In integrated mission mode, **1** creates a classifier, **2** creates a classifier successor, **7** opens the World Engine, and **9** opens compositions. The ordinary wizard remains available, with **0** entering mission mode and **8** reaching the Academy. Menu numbers are not maturity or authority levels. [3, 5]

5 Your first session: a complete synthetic cycle

Finish line: a native signed candidate, its readable result, a retained mission dossier and a byte-identical restored artifact. A negative local evaluation is still a successfully recorded lesson.

5.1 The optional guide companion

From the extracted **guide** directory, not the product directory:

```
| python3 START_LESSON.py
```

The companion asks for a language, the extracted product folder, a **new** practice folder and one of `classify`, `world` or `compose`. Choose `classify` for the simplest first lesson. Choose `compose` when your immediate goal is understanding predictor–planner records and shared resources.

The companion first verifies the product. It invokes the product’s explicit synthetic preparer and writes a populated teaching dossier **without approval**. It does not turn your company’s documents into data, silently choose independent proof or create a production permission.

Stage	Your deliberate action	Inspectable effect
Prepare	Choose a new folder and route	INPUT/PLAN. json and draft dossier, no training
Review	Read the synthetic scenario and omissions	Native settings and a HOLD approval state
Authorize local research	Type yes or oui, then name the local operator	A declared teaching-only approval, not authenticated corporate authority
Protect the key	Enter and repeat a hidden passphrase of at least ten characters	The companion passes it through a pipe, never an argument
Confirm	Inspect the actual native preflight , then confirm separately	A new exact native confirmation
Form and verify	Wait for the bounded run	RUN/, native result and signed artifact
Restore	Explicitly request the exercise	EXPORT/, RESTORED/, exact artifact-byte comparison

The companion’s `SESSION.json` records actual commands, exit codes and statuses. Native logs are saved beside it. Your passphrase is neither written into the session nor echoed into the logs. Evidence and model contents remain readable. A cancellation is recorded without pretending that a result was completed.

No continuation of training is implied. Reopening files lets you inspect a completed session; it does not resume an interrupted optimizer. To repeat a lesson, choose a new directory. To continue a real lineage, use the native successor workflow with an explicit evidence policy.

5.2 The entirely native route

The companion is optional. Run node `07_NATIVE/start.mjs` mission wizard from the product and enter all material fields yourself. It collects native settings first, then the mission dossier, then preflight and confirmation. It never treats a documentary HOLD as permission to train. The prompt-by-prompt cards in `03_COMMANDS/` identify the decision being made at each stage.

5.3 Read the first result

Open RUN/report.html, RUN/result.json and RUN/mission-dossier.html as local files. When a browser policy blocks navigation, use the JSON/text records and native verification. Do not weaken the policy. The classifier file is mission.successor.json, the World Engine file is world.successor.json, and the composition file is composition.successor.json.

Read, in order: **evidence origin** → **selected comparisons** → **local outcome** → **limitations** → **restore status** → **next decision**. Do not infer operating value from a completed training command. Public teaching keys and signed QA artifacts included with this guide are not identities to use for customer work.



Constitute the mission

Conceptual editorial fiction. Not a product interface, measurement, proof or permission.

PART / PARTIE II

Choose what deserves to be learned

6 Configure a real mission before supplying real data

The **SEIZE card** now provides the upstream decision: retain, repair, rent, form, partner, acquire, hedge, reserve, retire or prohibit. The **Mission Gym contract** describes the environment and omissions that make a comparison meaningful.

A useful mission names a repeated decision, not an aspirational category. “Improve procurement” is too broad. “Classify already identified supplier exceptions for human review using a documented six-signal mapping” is a possible classifier research mission. The product does not itself extract invoice evidence into that mapping.


The white paper’s mission screen asks for observable outcomes, credible alternatives, rights, constraints, evaluation and an accountable sponsor. Its decision gates are separate capital and authority decisions. [2, pp. 7, 13–14, 44–46]

Field	Good working answer	Warning sign
Decision and unit	One exception, one regular episode, or one shared-resource proposal	“Manage the company”
Outcome and horizon	A measurable event within a stated period	A vague improvement with no observable result
Incumbent and Beta	Named processes, dated versions and realistic constraints	A weak historical process chosen to make the candidate win
Evidence	Known origin, use rights, event identities and availability times	A folder of documents with no valid mapping
Hard limits	Enforceable numeric restrictions plus separately reviewed prose	Safety promised only in a mission title
Exit condition	Failed comparison, missing evidence, changed scope or budget exhaustion	More learning automatically authorized after every failure
Accountable owner	A person or organization able to make the next decision	A generated agent name acting as its own authority

Use the included mission worksheet. Define baseline, target, horizon, unit, planning opportunity and constraints deliberately. Some classifier fields retain required positive ranges even in a nonmonetary tutorial. A positive *planning opportunity* field is a declared input, not observed revenue. In a nonmonetary mission it must not be interpreted as observed revenue.

Three different constraints: a prose statement is retained but marked `BOUND_NOT_AUTOMATICALLY_COMPILED`; a supported numeric shield is actually enforced; a real operational policy requires its own external controls. Do not substitute the first for the second or third.

INTENT → PLAN → RETAINED RECORD



Beyond this boundary: independent custody, proof and separate admission.

A complete document does not remove the boundary.

Figure 2. A source-derived pathway with an explicit local boundary. SEIZE worksheets prepare a decision; the existing 20.0.1 workbenches execute bounded local research. Independent proof and admission remain separate.

7 Before formation: use SEIZE to choose the next decision

Source-derived operating discipline; translated into guide worksheets and the native mission dossier, not a complete implementation of the architecture. The Navigator asks which succession path is worth underwriting. The Gym paper asks what environment could make that path testable. Together they add the missing question before the training command: **What is the smallest useful piece of evidence that could change our decision?** [6, pp. 20–23, 7, pp. 23–26]

The canonical progression is **Successor Manifold → Mission Advantage Gradient → SEIZE Underwriting → Bounded AGI Jobs → Specialist ASI → Mission-Sovereign Successor**. It is an institutional progression, not six installed 20.0.1 menus. Alpha is a residual advantage, not a seventh chronological stage. The latter two states need the sources' comparative proof and separate admission; the guide and its tools do not assign either state.

7.1 The five questions, exactly in their source order

SEIZE step	Ask the owner	Retain before proceeding
S — Surface the succession event	What changed, and which incumbent assumption may no longer hold?	Dated event, source, affected assumption, cost of waiting and cost of premature change.
E — Evaluate the successor frontier	Why build rather than retain, repair, rent, partner, acquire, hedge, reserve, retire or prohibit?	Distinct alternatives, serious baseline and an anti-recommendation for each.
I — Instantiate the succession constitution	What mission, rights, hard gates, evidence boundary and bounded budget are being accepted?	Owner, beneficiary, outcome, units, horizon, stop-loss, evidence plan and accountable roles.
Z — Zero in on decisive proof	Which smallest diagnostic could alter this choice?	Named evidence object, complete cost, decision rule and reason to stop.
E — Elevate one challenger	Which complete candidate is ready to be frozen and examined?	Exact composition, unresolved limits and a Promotion Request—not a Release Certificate.

These step names and the anti-recommendation rule come from Navigator §9 and Gym §8.1. A reserve can be the right outcome. So can a recommendation to keep the incumbent. An attractive mission does not compel training. [6, pp. 21–22, 7, p. 24]

First practical use. Open 04_PRACTICE/MISSION_KIT/WORKSHEETS_EN.md. Complete the SEIZE card using one real decision, but no confidential raw records. Keep claims such as “a cheaper competitor is available” linked to dated evidence, not an unverified assertion. The source’s Successor Book can then hold an immediate option, challenger, reserve or hedge; 20.0.1 does not supply a live market radar that authenticates those options.

7.2 Buy information, not the appearance of progress

The sources define Value of Proof as

$$\text{VoP}(e) = \mathbb{E}_e \left[\max_a \mathbb{E}[U(a) \mid e] \right] - \max_a \mathbb{E}[U(a)] - C(e).$$

Here the evidence changes a decision before full commitment. 04_PRACTICE/ANALYTICAL/RESULTS.j son and the accompanying calculator implement a small, finite-state **teaching calculation**, not a probability estimator. All probabilities and utilities are declared assumptions. A test costs 5 utility units; its expected decision improvement is 18.8, leaving a net assumed value of 13.8. The untested build option has expected value –26, so the initial best option is retention at zero. The example shows why a diagnostic may be useful even when immediate construction is not. [6, pp. 10–11, 15–16, 7, p. 24]

A positive calculation does not override rights, safety or evidence-access restrictions. Unknown probabilities remain unknown. Do not back-solve a flattering probability from the budget you wish to spend.

8 Constitute the Mission Gym without confusing it with a world model

The Gym is the controlled representation against which a candidate is formed or assessed. A learned world program is the candidate’s **hypothesis** about dynamics. A policy chooses actions. The complete successor additionally contains evidence, tools, people, operations, rollback and authority. These are different objects. The Gym paper makes their separation the foundation of its triad. [7, pp. 3–5, 10–12]

For a first mission, write four sentences: **what is observed; which decision is made; what outcome later becomes available; and what would make the result unacceptable**. Then specify the complete environment. Do not substitute an impressive interface for a defensible observation contract.

8.1 The environment contract

Contract element	Concrete operator question
State and observation	What can actually be known at decision time? What is hidden, delayed or only estimated?
Actions and effects	Which labels are proposed, which actions were executed, and which effects are forbidden?
Transition or outcome	What generates the next observation: recorded reality, a fixed simulator, a learned hypothesis, or a human process?
Reward and hard gates	What outcome matters, and which breach cannot be compensated by another good score?

Contract element	Concrete operator question
Scenario and partition	Which episodes, entities and periods serve formation, challenge and transfer? What was already exposed?
Termination and truncation	Did the mission genuinely finish or fail, or did time, evidence, budget or authority run out?
Custodian and receipts	Who controls the environment, scorer, protected cases and exact version? What is actually preserved?
Rights and basis risk	What may be used or exported, and how can this environment misrepresent reality?

The source formalizes this as a constrained partially observable stochastic game. 20.0.1's bounded adapters do **not** implement arbitrary hidden-state games, strategic participants or live shadow interfaces. The field guide translates the specification into questions; it does not compile every answer into software. [7, pp. 5–9]

8.2 Five environment families, not five names for the same holdout

Source family	Intended role	What a 20.0.1 user can honestly retain now
Formation Gym	Repeated candidate construction.	Local learning data, explicit synthetic examples and formation records.
Adversarial Gym	Expose exploits, constraints and failure.	Deliberate local negative tests, labeled with their actual scope.
Transfer Gym	Challenge population, regime and failure-cluster changes.	Separately prepared eligible data and the limits of any local comparison.
Fresh-Proof Gym	Examine one frozen candidate outside formation access.	A proof request and proposed custodian requirements; no self-issued independent receipt.
Canary and Requalification Gym	Connect evaluation to controlled operation and renewal.	A manual readiness and renewal plan; no production tool or new authority.

The source's G0–G7 maturity ladder runs from narrative mission to replay, interactive environment, represented multiplayer institution, protected proof plane, canary, sovereign Gym asset and recursive succession environment. 04_PRACTICE/MISSION_KIT/GYM_MATURITY.md supplies the exact source stages with evidence questions. It deliberately assigns **no automatic G-level to 20.0.1 or a worksheet**. Showing roles does not establish independent custody. [7, pp. 8–9]

8.3 Keep the real-world gap visible

Gym basis risk includes missing mechanisms, absent rare events, inaccurate simulated humans, delayed outcomes, rights restrictions and distribution shift. The sources propose a basis-risk reserve; they do not supply the real gap for your company. Record each omission, the affected claim, the person responsible, a possible measurement and the condition requiring hold or narrowing. A numeric reserve is not permission to ignore a failed hard gate. [7, pp. 39–40]

For an invoice mission, a perfect synthetic ledger does not establish messy-document extraction, causal recovery of cash, fraud identification or a right to change payments. For 20.0.1 composition, normalized heat and fouling do not establish calibrated factory temperatures or physical safety. This is why the environment needs its own version and review, independently of the candidate's version.

9 Bind intention to the exact executable plan

SEIZE helps decide whether research is worthwhile. The **native mission dossier** connects that reasoning to the configuration the product will actually execute. The dossier does not replace the native plan.

The contract is `successor-omega/mission-dossier/v1`. It binds a route, mission identity, title, owner, objective, evidence policy, decision unit and declared bounded mapping. It records the succession event, alternatives and anti-recommendations, next decisive evidence, stop condition, environment gaps, rights reference and renewal requirements. [3]

9.1 A useful example of a narrow mission

Fictional teaching organization: an industrial operator wants to prioritize already structured exception records for human review. Its mission is not “run the factory.” It is “classify these normalized cases under the declared six-signal mapping, compare with the supplied incumbent and Beta snapshots, and retain the result for review.”

The owner keeps the alternatives explicit. Retaining the current process may avoid integration cost. Repairing that process may be cheaper than training. A candidate is worth considering only when the next evidence can change the decision. A larger budget does not remove the need for that test.

The mapping may be UNKNOWN, UNSUPPORTED or DECLARED_BOUNDED_FIT. The last value is an operator declaration, not an external engineering finding. Approval is initially HOLD. Only an appropriately reviewed LOCAL_RESEARCH dossier for a form/repair option can proceed to native preflight.

9.2 The exact commitment

The binding records both the original UTF-8 bytes and their parsed semantic identity. It references the native plan **excluding only its confirmation and the binding itself**. Changing an evidence path, comparator, objective, economics or workspace changes the bound plan. Changing the original dossier bytes, even after an apparently harmless edit, requires review and a new binding.

The bound dossier travels inside the signed record. Reopening a completed artifact does not need its former absolute dossier path. That does not mean a relocated *new training plan* may be used without checking its paths and creating a fresh confirmation.

```
node 07_NATIVE/start.mjs mission template \
  --plan ../manual-mission/PLAN.json \
  --out ../manual-mission/DOSSIER.json
node 07_NATIVE/start.mjs mission review \
  --file ../manual-mission/DOSSIER.json
node 07_NATIVE/start.mjs mission bind \
  --file ../manual-mission/DOSSIER.json \
  --plan ../manual-mission/PLAN.json \
  --out ../manual-mission/BOUND.json
```

These commands use a separate **manual-mission** folder created by the direct native preparer. The teaching companion instead keeps its plan in **INPUT/**. Do not combine the paths of those two examples. Complete the draft before binding. Keep BOUND.json and CONFIRMED.json beside the original plan so relative evidence paths retain their meaning. A new output name is required: the tool does not silently overwrite an existing record.

A READY_FOR_LOCAL_NATIVE_PREFLIGHT review means only that the required declarations are present and internally acceptable to this schema. Rights are not authenticated. The broader architecture’s independent verifier, protected proof and Authority Envelope remain separate institutions.

10 Turn the mission into twenty-one bounded work-and-proof jobs

The following job names preserve **Gym Appendix C**, printed pp. 61–62. The short outputs are operator paraphrases. The checklist is included as JSON and bilingual Markdown with owner, reviewer, status, artifact reference, budget and next decision. Every initial status is UNASSESSED: a generated row is not completed work. [7, pp. 24–26, 61–62]

No.	Source job	Required retained object
1	Mission Constitution	Frozen purpose, beneficiary, constraints and prohibited actions.
2	Incumbent Baseline	Current performance, complete burden and operating state.
3	Alternative Set	Serious frontier, specialist, human, hybrid and no-action choices.
4	Rights and Data Constitution	Source, licence, confidentiality, retention and reuse record.
5	Mission Gym Constitution	State, observations, actions, outcomes, gates and basis-risk plan.
6	Benchmark Constitution	Representative work, scorer, critical failures and fixed hurdle.
7	Oracle and Teacher Study	Declared teaching capability and prohibition of hidden calls.
8	Failure Taxonomy	Severity, causal hypotheses, abstention and stop conditions.
9	Architecture Frontier	Diverse complete candidates and reserve options.
10	Data and Experience Formation	Eligible examples, simulations and quality controls.
11	Specialist Model Formation	Exact adaptation, fitting or algorithmic construction.
12	Tool and Retrieval Formation	Tools, checks, planners, interfaces and their own evidence.
13	Verifier Formation	Evaluation method, calibration, attacks and protected scoring.
14	Security and Privacy Proof	Threat model, boundaries, credentials and incident controls.
15	Complete Economics	Formation, operation, review, maintenance and unwind burden.
16	Human Burden and Agency	Escalation, contestation, supervision and accountable judgment.
17	Resilience and Rollback	Known-good restoration, fallback and continuity.
18	Protected Fresh Evaluation	Frozen challenger, unseen work and complete denominator.
19	Independent Validation and Acceptance	Separately governed review of evidence and fitness.
20	Chronicle and Authority Admission	Scoped admission, expiry, revocation and negative history.
21	Capital-to-Capacity and Requalification	Realized-value allocation and the next fresh comparison.

Use the portfolio proportionately. One person can draft several local research records; that does not make their review independent. 20.0.1 can contribute bounded training, local comparison, exact records and restoration to several jobs. It does not complete jobs 13–21 merely by producing similarly named files. The full source-to-20.0.1 map distinguishes native support, guide-side preparation and external institutional work.

Keep the other source's portfolio identifiable. Navigator §12.6 also describes a 21-job formation portfolio. Gym inserts **Mission Gym Constitution** at job 5 and groups **Independent Validation and Acceptance** under job 19; Navigator has separate validation and acceptance jobs 18 and 19. 04_PRACTICE/MISSION_KIT/JOB_CONCORDANCE.json preserves that numbering and wording without merging the underlying accountable decisions. The checklist above is explicitly the Gym version. [6, pp. 29–30, 7, pp. 61–62] A work-and-proof job is not an automatically dispatched agent task, blockchain job or paid transaction.

11 Preserve the distinct protocols, symbols and decisions

The two attached publications extend the foundation of the user guide, but they are not runtime patches. Use a source-qualified protocol name before copying an equation or interpreting a receipt. Do not choose a definition after seeing which one produces a favorable result.

11.1 Three related Alpha accounts are not one interchangeable formula

The Navigator's compact account is

$$\alpha_{\text{NAV}} = U(S) - \max\{U(I), U(B)\} - C_m - R_m - D_m,$$

where C_m covers construction and lifecycle cost, R_m risk including tail and authority exposure, and D_m unresolved Proof Debt and dependency burden. [6, pp. 14–15]

The Gym account adds an explicit environment-to-reality reserve:

$$\alpha_{\text{GYM}} = U(S) - \max\{U(I), U(B)\} - C_m - R_m - D_m - \hat{e}_{\text{gym}}.$$

The reserve is not estimated merely by having a simulator. [7, pp. 21–22, 39–40]

The later white paper uses the lower-bound comparison and deductions retained by the 20.0.1 paper-aligned protocol:

$$\alpha_{\text{WP}} = \text{LCB}[U(c) - U(b^*)] - C - R - H - I - D.$$

Here D means **drift and impairment**, not the Navigator's Proof Debt/dependency aggregate. H is human burden and I imitation/commoditization. The symbols must not be aliased by spelling alone. Existing signed results keep their original semantics. [2, pp. 30–32]

Operational rule. Use unique burden IDs. For every item, record units, horizon, amount or unknown status, whether it is already inside utility, and exactly one official deduction. The 04_PRACTICE/MISSION_KIT/BURDEN_REGISTER worksheet provides explicit mappings into the selected account. A crosswalk is a proposed mapping to be approved before use; it does not rewrite any source or calculate an unmeasured reserve. Qualitative Proof Debt stays qualitative until a declared conversion is justified. No sums are calculated from blank amounts.

11.2 A fixed baseline differs from a hindsight casewise envelope

The Gym's main reference equation selects the strongest reference by mission utility. Its Appendix D.5 illustrative pseudocode instead subtracts $\max(\text{incumbent.utility}, \text{alternative.utility})$ inside each case. These need not estimate the same quantity. The latter is a casewise envelope that may not be implementable without outcome hindsight. The guide preserves both locations rather than calling them equivalent. [7, pp. 21, 63]

For two equally weighted hypothetical cases, incumbent utilities (10, 0) and Beta utilities (0, 10) both average 5. A candidate at (6, 6) beats either fixed baseline by 1, but loses 4 against the per-case envelope. This exact calculation is included as a teaching test, not an observed product result. Precommit which estimand is official. Do not change the 20.0.1 evaluator or a historical receipt to resolve the source difference silently.

11.3 Keep three decision frameworks distinct

SEIZE structures an underwriting microcycle. **Underwrite** → **Execute** → **Accept** → **Admit** → **Allocate** assigns five different accountable decisions in the attached sources. **D0–D4** in the later white paper is a separate executive sequence: Underwrite, Build, Prove, Admit, Renew/retire. Related purposes do not make these labels interchangeable software states. [6, pp. 20–24, 7, pp. 29–32, 2, p. 7]

The Navigator says *Singularity Exposure Ratio*, whereas the Gym says *Succession Exposure Ratio* for L_a/H_c . Both are preserved with source-qualified names. No adaptation lag or capability half-life is measured by the guide. Similarly, manifold, gradient, least-action and gauge-invariance language is identified in the sources as a design analogy. 20.0.1 does not infer a calibrated institutional metric tensor or optimize that whole architectural geometry. An expression using G^{-1} requires a suitable invertible metric; a declared semidefinite form alone does not supply that inverse. [6, pp. 1, 5–7, 7, pp. 2, 14–16]

11.4 Access, commercial offers and authority remain separate

Navigator §16 describes two eligibility routes for **its public browser-local institution**: direct qualifying AGI Club ownership or a direct qualifying AGIALPHA holding. Its access receipt explicitly creates no authority. Those historical scope-specific terms are not added to the separately licensed offline 20.0.1 product. This integration introduces no wallet connection or balance check. [6, pp. 36–37]

The attached commercial ladders contain illustrative engagement prices and proposed services. They do not amend 20.0.1's licence, price, support contract or customer rights. An architectural reference API in the Gym appendix is pseudocode, not a command newly installed by attaching its PDF. [7, pp. 52–55, 62–64]



Make hypotheses executable

Conceptual editorial fiction. Not a product interface, measurement, proof or permission.

PART / PARTIE III

Learn, compare, and use a candidate

12 Workbench A: customer-specific classification

Use this route for bounded tabular decisions mapped to the six governance signals. It fits small neural candidates and compares them with explicitly constituted alternatives. It does not read arbitrary documents, train a general-purpose agent or identify the causal return of every possible action. [3]

12.1 Data you must prepare

A cases-v1 document uses schema successor-omega/cases/v1 and a cases array. Each case has an id, six observation values, preferredAction and positive valueWeight. Optional fields include group, numeric time, declared role, incumbentAction and betaAction. Supported labels are observe, recommend and reversible_experiment.

The legacy eight-column CSV is:

```
evidence_strength,outcome_gap,intervention_fit,stakeholder_risk,  
reversibility,drift_risk,preferred_action,value_weight
```

The ten-column form adds incumbent_action, beta_action. The actual header is one line; the line break above is for reading only. Use 30–1,000 cases with at least ten training and five development cases. Do not truncate an oversized dataset silently. Without explicit event IDs, identical observations can be ambiguous. Prefer the identified JSON adapter for real history and explicit role control.

Separate snapshot files are not joined by accidental row order. Prepare an audited, case-ID-aligned document. Missing comparator actions are not invented, and unsupported actions fail validation. The transformation from business evidence into the six signals remains customer work.

12.2 Comparator choices that really change evaluation

Choose the incumbent independently: local_rule or snapshot. Choose Beta: local_proxy, snapshot or absent. Imported Beta currently requires the imported incumbent coverage alongside it. A proxy is a local proxy: **no external model was evaluated**. A label such as a provider name does not authenticate that provider.

When a snapshot column exists but you do not select it, the product asks you to acknowledge exclusion. Preserve that decision. Do not silently ignore an unfavorable column. Constitute the label, version, rationale, coverage, available collection window and costs before examining results.

12.3 Complete the route

Choose menu 1. Edit every material mission field, choose the evidence path and format, confirm rights, choose the partition policy and comparators, then set the accounting method and evaluation rule. Read preflight, confirm, choose a new output directory and enter your secret through the hidden prompt.

The run writes `mission.successor.json`, `result.json`, `evaluation.json`, `chronicle.json`, `proof-request.json`, `report.html` and a readable note. Inspect the per-case ledger—not merely the aggregate label. An imported comparator’s action must be visible as the action that was actually evaluated.

The supported numerical shield forces observe when evidence strength is below 0.34 or drift risk is at least 0.82. It downgrades `reversible_experiment` when stakeholder risk exceeds 0.64 or reversibility is below 0.48. These are implemented adapter rules, not a custom safety theorem. The guided display exposes its evidence, drift, risk and reversibility thresholds. Candidate selection uses the declared development rule; evaluation follows freezing. A good training score is not a holdout result.

13 Read comparisons without fooling yourself

The **source-protocol concordance** distinguishes the Navigator, Gym and later white-paper accounts; in particular, their symbol D does not always mean the same burden.

The classifier supports a nonmonetary protocol and a separately versioned paper-aligned scenario protocol. Historical artifacts retain their older arithmetic. Do not recompute an old signed result under a new definition and call it the same receipt.

For the paper’s framework, the gross-utility baseline is

$$b^* = \arg \max_{b \in \{\text{incumbent}, \text{Beta}\}} U(m, b),$$

and the conceptual adjusted advantage is

$$\text{Alpha}(m, c) = \text{LCB}[U(m, c) - U(m, b^*)] - C - R - H - I - D.$$

The deductions are incremental operating/integration cost, risk reserve, human burden, imitation/commoditization reserve and drift/impairment reserve. They must use the same units and horizon and be charged once. This is a mission governance template, not a universal accounting standard. Independent proof requires more than this arithmetic. [2, pp. 30–32]

Gross versus net ranking matters. A baseline with larger gross utility can have higher costs. A precommitted gross-ranking protocol must not change to net ranking after you see a favorable result. Keep the official definition fixed and label any alternative diagnostics separately.

Missing is not zero. An absent Beta disables Beta-relative conclusions. Unknown financial costs must not be entered as zero. Missing required monetary inputs block that evaluation; the nonmonetary route remains available. A nonmonetary comparison is preferable to invented dollars. A synthetic payoff table describes scenario preferences; action snapshots alone do not identify the outcomes of unexecuted actions.

Uncertainty is method-specific. A sensitivity allowance is not a general confidence guarantee. A bootstrap interval depends on its sampling unit and assumptions. Repeatedly adapting to the same

local holdout does not create fresh evidence. Preserve failures, abstentions and incomplete coverage in the denominator.

13.1 Read the result in a fixed order

Field or evidence	Read it as	Do not infer
execution: COMPLETE	The run finished and wrote its result.	The mission succeeded.
localOutcome: PASS	This local protocol's declared comparison passed.	Real savings or independent proof.
LOCAL_PREDICTIVE_EVALUATION	Prediction was evaluated on eligible recorded episodes.	Better operating decisions or financial gain.
INSUFFICIENT_ELIGIBLE_EVIDENCE	The lineage cannot support the requested evaluation.	Permission to repartition exposed cases as fresh.
INFORMATIONAL_PROPOSAL	The operator path produced a hypothetical proposal.	An approved service order.
preferred or preference	A route-specific local selection state.	Production-serving admission.
independentProof, productionAdmission, realizedAlpha	In these lessons: not established, not admitted, zero.	An earned designation from a file suffix.

The exact key names differ by artifact schema; use that route's report and reader. A saved report describes formation history. For a composition, `compose` status also consults the current clock and operator journal. The guide does not collapse these states into one "success" badge.

14 Workbench B: executable predictive worlds

Choose `world wizard` when action consequences unfold over time and the supported temporal representation fits. The engine searches a bounded typed grammar, fits parameters from training episodes, selects with development evidence, calibrates separately and freezes before eligible evaluation. A neural proposal ranker can order search; it is not a frontier language model or a guarantee of better final search quality. [3]

```
| node 07_NATIVE/start.mjs world wizard
```

The built-in mission packs are `reliability` and `workload`. The wizard exposes operation, mission, pack, evidence policy, search budget and planner settings. Select a supplied example explicitly; do not pass a classifier CSV as temporal evidence.

14.1 The temporal contract

The trajectory schema carries source and rights, episode ID, group, role, initial history and ordered steps. Each step records observation, availability and action times, outcome availability, unit step duration, current observations and inputs, executed action, next observation, outcome and terminal status. Behavior probabilities are retained when meaningful; they are not invented.

Roles are training, development, calibration and evaluation. The supported format contains 1–500 complete episodes, 3–64 steps per episode and at most 20,000 transitions. The representation must have enough eligible role counts for the chosen fitting, calibration and evaluation rules. The syntactic maximum is not a sufficient evidence sample.

Irregular intervals, censored outcomes, arbitrary missing sensors and general hidden-state estimation are not silently repaired. The actor receives only current permitted inputs and bounded history. Future outcomes, latent simulator state and regime labels must not enter the actor's features.

14.2 What the World Engine actually outputs

A typed program is a testable hypothesis about next state. It can contain arithmetic, products, conditions and eligible history dependencies. A structural revision changes represented dependencies or conditional mechanisms; a new coefficient alone is not the same claim.

Inspect `world.successor.json`, `result.json`, the program, candidate archive, contradiction and structural-change records, evaluation and `proof-request.json`. Use the record's own links; do not edit signed JSON to make its story more convenient.

```
node 07_NATIVE/start.mjs world verify --file WORLD_RUN/world.successor.json
node 07_NATIVE/start.mjs world predict --candidate WORLD_RUN/world.successor.json \
  --context CONTEXT.json --actions 0,1,0
node 07_NATIVE/start.mjs world plan --candidate WORLD_RUN/world.successor.json \
  --context CONTEXT.json
```

These commands compute model-dependent scenarios. They do not run the proposed actions in reality. The prediction sequence must use actions permitted by that specific program.

15 Interpret uncertainty, contradictions and revision

A one-step calibration radius describes the implemented residual method under its assumptions. It is not a six-step safety guarantee, a causal confidence statement or permission to cross a resource limit. Free-running rollouts propagate generated state; teacher-forced accuracy on actual future states answers a different question.

A residual event says a prediction and observation disagree. Possible causes include model error, sensor error, timing, changed dynamics or inadequate data. Do not turn one surprising noisy observation into proof of a unique mechanism.

Use a separate revision branch. Preserve the original program and the observation that prompted review. If previously evaluated evidence is used to guide a repair, record its new formation role and obtain new eligible evaluation evidence for the later claim.

Outcomes to recognize: PROPOSAL is a hypothetical plan; ABSTAIN_INSUFFICIENT_CALIBRATION asks for adequate calibration; ABSTAIN_UNCERTAINTY says the declared radius is too large; BUDGET_EXHAUSTED is not a completed best plan. REGRESSION_REQUIRES_SCOPE_REVIEW means a child improved something while degrading retained behavior and needs an explicit scope decision.

For cancellation, Ctrl+C requests the bounded worker to stop. The checkpoint is not a completed qualification. `world resume` replays formation under a newly confirmed plan and bounded budget; it does not restore an invisible partially approved model or make used evidence fresh. Preserve the interrupted record.

16 Workbench C: compose intelligence under shared constraints

A predictor alone does not settle how several assets should share electricity, material or a maintenance team. This workbench adds a complete, inspectable planning composition around a compatible predictor.

```
| node 07_NATIVE/start.mjs compose wizard
```

16.1 Choose the predictor honestly

train-neural fits a local multilayer network and selects on development prediction error. import-world verifies a compatible signed world and observation contract. import-neural imports a bounded model and permitted calibration evidence, then recalculates calibration. Import is not fresh training, and a source label is not provider authentication.

A composition's native evaluation concerns recorded prediction performance. CHILD_BETTER_PREDICTION is not evidence of greater plant contribution. To compare operating policies, a separately authorized evaluator needs actual outcomes from each legitimate policy trajectory or justified prospective evidence. The guide does not infer money from prediction error.

16.2 Constitute the resource contract

The importable resource-contract JSON contains exactly assets, objective and scope. The supported adapter has one to six process cells. Each observes two normalized states and two inputs with two history entries; actions are 0 idle, 1 produce, 2 service.

The objective contains scenario prices, a yield function, terminal heat/fouling charges, a heat threshold and inventory capacity. The scope contains validity times, population, observation-age limit and intervals for feed, electricity, crew, tariff and observed delivery. These are material configuration—not defaults to accept merely because the screen is complete.


Supported scenario price fields do not automatically become measured financial outcomes. Text stays BOUND_NOT_AUTOMATICALLY_COMPILED. A real physical mapping needs sensor definitions, units, instruments and independent engineering evidence.

16.3 Choose a planning method

Method	Exact interpretation	Good reason to inspect it
ratio	Bounded greedy heuristic using incremental value and normalized electricity	Low-compute coordination; inspect priority and normalization assumptions.
exact-local	Exact current joint action over a surrogate of local continuations	A reference for current shared feasibility, not exact joint future control.
coupled-beam	Bounded search carrying resource constraints through hypothetical future stages	Examine delayed competition; do not call it a global optimum.

Horizon is one to three steps. Beam width is one to sixteen. Prediction and joint-inspection budgets are bounded. ZERO assumes no future delivery; HOLD_OBSERVED repeats the observed delivery hypothetically. Neither reads tomorrow's delivery. More compute or a longer horizon does not guarantee better decisions in the true environment.

A SIGNATURE IS NOT A PERMISSION



These checks do not establish physical truth or independent proof.

Figure 3. The public composition path checks separate records; a signature does not establish truth.

17 Pin, inspect and propose

The signed artifact contains a self-issued key identity. The trust file represents a separate local operator decision. Keep them conceptually and operationally distinct.

```
node 07_NATIVE/start.mjs compose verify --file RUN/composition.successor.json
node 07_NATIVE/start.mjs compose pin --file RUN/composition.successor.json \
  --trust operator-trust.json --expected FULL_DIGEST --actor "Mission owner" \
  --reason "Reviewed for local research"
node 07_NATIVE/start.mjs compose plan --file RUN/composition.successor.json \
  --trust operator-trust.json --context CONTEXT.json --out PROPOSAL.json
```

Replace FULL_DIGEST with the entire identity returned by verification of the inspected file. Do not pin an identity blindly from an untrusted message. Do not include literal placeholder words in your command. Use a new output path for every proposal.

The context declares population and UTC observation time, each cell's current states, inputs and history, available feed, energy, crew, tariff and observed delivery. Cell IDs must match the assets. Availability is physical input, not a forecasted resource you hope to receive.

The public planning path checks identity, revocation, validity, observation age, population, resources and uncertainty before computation. `compose status` combines frozen formation history with current operator/time state; a saved report remains historical.

```
node 07_NATIVE/start.mjs compose status --file RUN/composition.successor.json \
  --trust operator-trust.json
```

The trust journal does not solve a fully compromised operator, replacement of the trusted journal, or fabricated physical provenance. A clean signature and a current date do not make an incorrect sensor value true. Pinning is necessary for the supported path, not sufficient for independent proof.

18 Compare compositions and decide locally

Use `compose decision` with `retain`, `select`, `archive`, `impair` or `revoke`, a named actor and a reason. These choices preserve history. A child never silently replaces its parent. Rehearsal and insufficient-calibration compositions cannot become preferred.

```
node 07_NATIVE/start.mjs compose decision --file RUN/composition.successor.json \
  --trust operator-trust.json --decision retain --actor "Mission owner" \
  --reason "Keep this candidate while the child is evaluated"
```

RETAIN THE PARENT, TEST THE CHILD



Compare frozen generations

Retain / select locally / archive — explicit decision

No inherited proof or authority.

Figure 4. The retained local candidate is not mutated while the learning branch is formed.

retain is not a new superiority finding. select is a local preference, not a production-serving admission. archive preserves a nonpreferred alternative. impair records why applicability needs review. revoke stops local proposals under that trust journal and clears its preference without deleting the object.

The expert compose compare command is a **behavior diagnostic** on supplied contexts. It reports whether actions change. It does not identify the decision value of alternative trajectories and does not replace the pinned operator planning path.

```
node 07_NATIVE/start.mjs compose compare --parent PARENT.json --child CHILD.json \
  --contexts CONTEXTS.json --out BEHAVIOR.json
```

The contexts file is an array, limited to the supported count. Assets and objectives must be compatible. A changed planner, objective or population can make native succession incomparable; branch explicitly rather than declaring improvement from unrelated scores.

19 Create a successor deliberately

Before creating a child, verify the predecessor and inspect its mission, adapter, evidence origin, roles, exposure, comparators and local claim state. Copying an artifact or inheriting a mission title is not inherited trust.

Evidence policy	What it means	What it does not mean
new	Newly supplied permitted evidence; parent retained as history	Automatically authentic, independent or never previously used
reuse	Explicit reuse of eligible parent evidence and exposure history	A fresh holdout
append	Eligible historical and new cases with identity/partition rules	Blind concatenation followed by a new random split
rehearsal	Explicit synthetic practice branch	A customer deployment candidate by default

A blank policy stops. An unreadable file does not become rehearsal. An already inspected evaluation event cannot become fresh by renaming the file, editing the label, or changing line endings. With explicit authorization it can be retired into training, but its recorded exposure remains.

The product distinguishes a parent-key-authorized continuation from a referenced_fork. A self-declared organization name is not authorization. Use the appropriate parent-secret route when required; never place private keys in lineage metadata.

The exact frozen parent and child must be compared under compatible mission semantics on eligible common evidence. A tied, inferior, insufficient or incompatible child is a legitimate result. The current local preference does not change until you explicitly decide. Keep a better-supported predecessor.

Practical improvement strategy: change one identifiable component or one evidence policy at a time. State the hypothesis before training. Retain a baseline run. A new evidence cohort and a new planner in the same step make attribution harder even when the complete composition is useful.

SUCCESSOR Ω / USER GUIDE Ω

Retain the institution



Conceptual editorial fiction. Not a product interface, measurement, proof or permission.

PART / PARTIE IV

Retain value without losing control

20 Restore an exact record without inventing permission

A restoration is successful when the exact permitted record travels and can be inspected under the relevant reader. It is not successful merely because a folder was copied.

20.1 What the key protects

Only the private signing key is encrypted. The mission, source evidence, learned parameters, dossier, provenance and most metadata are plaintext. Encrypt storage separately when required by your organization’s controls. Avoid uploading real evidence or secrets to a support channel merely because the file extension looks like an opaque container.

A public signature check can succeed without your passphrase. Possessing the private key is a separate test; parent-authorized continuation is a further explicit use of it. Losing a passphrase does not rewrite the public result, and a public result does not recover a lost private key.

20.2 Match the export family to the artifact

Artifact	Export/restore command prefix	What does not transfer automatically
mission.successor.json	none	Preferred local candidate, authority
world.successor.json	world	Local preference, independent proof
composition.successor.json	compose	Operator trust pin, preference, claim currency

For a composition:

```
node 07_NATIVE/start.mjs compose export \  
  --file ../first-mission/RUN/composition.successor.json \  
  --out ../first-mission/EXPORT  
node 07_NATIVE/start.mjs compose restore \  
  --dir ../first-mission/EXPORT \  
  --out ../first-mission/RESTORED  
node 07_NATIVE/start.mjs compose verify \  
  --file ../first-mission/RESTORED/composition.successor.json
```

Keep **every file the native export creates**. Version 20.0.1 validates exact directory members, complete unique checksum coverage, regular-file requirements and metadata consistency before the restoration write. Repeated checksum lines do not substitute for missing entries. Symbolic links are not a portable-export shortcut.

20.3 Effective restrictions survive composition and succession

A permissive new wrapper does not erase a restrictive source. A restricted mission dossier or known ancestor restriction can prohibit full export. Importing a signed world into a composition retains an authenticated effective-export decision distinct from the original declared source rights.

Historical objects remain inspectable through historical readers. An older imported composition lacking required source-export information may be refused for full export. The correction is not to erase provenance: recover the exact original world, review its permissions, and form a new authorized composition. [3, 5]

These are controls over **known declarations and recorded history**. They do not prove legal title, perfect ancestry or global non-use. Follow the source paper's distinction among owned, licensed, portable, replaceable and excluded components. [2, pp. 36–37]

20.4 A restoration drill you can complete

Retain the original artifact and export. Restore into a new directory. Compare the signed artifact bytes. Reopen it with the native reader. Inspect the mission, negative results and rights. For compositions, inspect the current clock and scope, then make a **new explicit pin** before any hypothetical proposal. Record what was tested and what remains unknown. Do not claim an externally operated supplier-substitution test from this same-machine drill.

21 Troubleshooting without weakening the boundary

What you see	Likely meaning	Correct next step
node not found	Runtime is absent or not on the terminal path	Install an approved supported runtime; open a new terminal; check <code>node --version</code> .
Cannot find <code>07_NATIVE/start.mjs</code>	Wrong current folder or archive not extracted	Navigate to the customer root; do not invent a different script path.
Permission denied on a launcher	Extraction did not preserve executable mode	Use <code>node 07_NATIVE/start.mjs</code> wizard directly; do not run as administrator merely to bypass it.
UNKNOWN_OR_DUPLICATE_FLAG	Wrong command family, obsolete example or duplicated flag	Read current help / compose help; <code>selftest</code> has no hyphen.
REQUIRED_ANSWER / missing evidence choice	A required question was left blank	Restart deliberately; absence is not a request for demonstration data.
Confirmation mismatch	Inputs, plan or exposure state changed	Inspect the difference and run a fresh preflight; do not paste an old token.
OUTPUT_EXISTS or file exists	Destination would overwrite retained work	Choose a new destination. Preserve the prior attempt.
SECRET_LENGTH / BAD_SECRET_OR_VAULT	Unsupported secret length, wrong secret or damaged vault	Use the correct secret and original file; do not edit the signed payload.
STALE_OBSERVATION / future observation	Timestamp is outside the contract	Supply genuinely current authorized evidence. Regenerate only explicitly synthetic demonstrations.
EXPIRED_OR_NOT_YET_APPLICABLE	Composition time window is not current	Reconstitute/requalify as needed; do not change the system clock to get a pass.

What you see	Likely meaning	Correct next step
OUTSIDE_DECLARED_SCOPE / population mismatch	Input or population differs from the declared envelope	Retain the refusal and form a separately qualified scope, not a wider unchecked label.
Insufficient eligible evaluation/calibration	Data roles or exposure cannot support the requested result	Collect eligible data or retain a trained-but-not-evaluated result.
Budget exhausted	No complete answer within the declared bound	Review complexity and budget in a new plan; a partial search is not an optimum.
Browser blocked by administrator	Host policy prohibits the navigation	Use the native console or an administrator-approved host. Do not disable controls.

For locks and .pending directories, first confirm that no process still owns the operation. Preserve failure records. Do not remove an active lock or call a partial artifact complete. The World Engine has a documented checkpoint-replay route; composition formation does not provide an equivalent resumable checkpoint command in this release. Use a new formation attempt after inspection.

Native compose compare is not the public applicability gate. Use compose plan with a trust journal for operator proposals. Comparing actions offline does not establish value or current authorization.

22 Academy, support and frequently confused features

The eight native Academy lessons cover mission, evidence, comparison, freeze, custody, authority, succession and the next human decision. Reading progress is saved and reopened; it is not a certificate. The larger Executive Academy and historical papers remain in the portable archives and publication bonus directory. [3]

20.0.1 language detail: menu 8 invokes the retained Academy endpoint, which defaults to French rather than inheriting the main menu's language. For an explicitly English native Academy session, the inspected and exercised route is:

```
| node 07_NATIVE/legacy-start.mjs academy --language en
```

Use --language fr for French. This is a specific retained Academy route, **not a recommendation to use legacy creation or succession commands**. In the Academy, Enter marks the displayed lesson read; q saves and exits. Retain the printed path to academy-progress.json and supply it at the next session to resume. The progress schema/release metadata may reflect its retained historical module.

Can I upload arbitrary PDFs to train a specialist? Not through these bounded native adapters. Map permitted content into the supported contract with documented transformations.

Does a new title change the numerical world? No. Unsupported prose stays bound but not automatically compiled.

Does a positive result make me money? No. It may justify another research decision; real outcomes, costs, uptake and authorization are separate.

Can I read an export on another computer? Often the file is readable, but the relevant runtime, verification and restoration checks still apply. Do not call one file copy independent portability qualification.

Does the product require a model service or wallet? The inspected local routes do not. Optional retained tools have their own documented dependencies. No external provider is enabled by this guide.

For launch support, consult the unchanged support policy in 04_LICENSE_AND_POLICIES_EN_FR/. It names info@quebec.ai, 30 calendar days from confirmed payment, Montréal business hours and a three-business-day initial-response target—not guaranteed resolution. This guide does not extend that contract or include confidential data preparation, external proof or production operations.

Send support the release ID, OS/runtime version, exact command, error code and a redacted reproduction. Do not email your passphrase, private keys, full confidential evidence or an unredacted operator journal unless separately authorized through an appropriate channel.

23 A practical operating cadence

This is a proposed cadence, not an automatic 20.0.1 monitoring service or a promised schedule.

First session: verify the package; complete a synthetic exercise; reopen and restore it; explain one refusal. Do not introduce confidential operational evidence until the file and key boundary is understood.

First real research cycle: write the mission contract; audit the evidence mapping and rights; constitute comparators; reserve evaluation roles; freeze the plan; train within budget; inspect adverse cases; retain or reject explicitly. End with a decision record, not a compulsory success claim.

At each subsequent research session: verify the exact artifact; inspect current time, scope, evidence changes and local journal; confirm the correct workspace; preserve the parent; authorize only a specific branch. If the data, costs, resource contract, predictor or planner changes materially, create a new identity and proportionate evaluation.

At review: compare the currently available alternative, actual human burden and restoration readiness. Treat expired, revoked or unsupported applicability as a reason to stop the associated proposal path. Preserve the old result as history. 20.0.1 has no background plant watcher that independently discovers every such change.

At transfer: verify rights, export, restore in a new location, check bounded behavior, then make a new trust decision. Do not distribute a customer's private evidence merely because the package can technically export it.

24 From a local specialist to an admitted institution

The Gym families and maturity questions and 21 bounded jobs make this broader path explicit without promoting local outputs into protected proof.

The architectural roadmap is larger than this release. It separates mission underwriting, formation, exact freeze, independent proof, accountable admission and renewal. Its independent-proof levels depend on actual control arrangements, not on how many keys, screens or subprocesses are involved. [2, pp. 27–33, 45–46]

Stage	Work you can retain	Evidence still needed for the stronger claim
Local formation	Plans, trained models, programs, local records and negative tests	Authentic customer outcomes and adequate comparison design
Frozen challenge	Exact composition and protocol	Appropriate independent protected-evidence custody
Independent proof	A separately governed proof/failure receipt, if actually obtained	Conflict disclosure, replication and applicability conditions

Stage	Work you can retain	Evidence still needed for the stronger claim
Production admission	A customer-granted envelope, if actually granted	Sector-specific controls, monitoring, rollback and accountable operators
Renewal	New evidence and a new explicit decision	Current baseline, unexpired scope and proportionate re-proof

20.0.1 supplies no button that awards the final two stages. A local trust pin is not an external authority envelope. A model may be genuinely trained and useful while remaining a candidate rather than an admitted Specialist ASI.

The long-horizon value proposition is **retained learning plus disciplined reinvestigation**: an operating mission identifies a resource bottleneck; a materials or manufacturing mission tests an improvement; actual surplus can fund deliverable capacity; the next mission inherits useful methods and failures. This guide does not demonstrate autonomous construction, unlimited compounding or stellar-scale power. Each physical and economic step has its own proof obligation.

25 Use SUCCESSOR Ω for the highest useful value

Recommended practice. Choose a mission whose value can survive a strong comparison, not a mission whose story permits the largest multiplier. The best initial research candidate combines repeated decisions, observable outcomes, costly errors, valid rights, a bounded adapter and a serious baseline. [2, pp. 44–46]

A practical portfolio begins with one mission and three budgets: a data-preparation budget, a formation/comparison budget and a separately authorized proof budget. Keep the spending decision with the human sponsor. The software's native research loops do not run an investment committee.

Decision card	Record before expanding
Observable advantage	Actual paired outcome definition; distinguish prediction, decision utility and money.
Competitive relevance	Incumbent and a credible current alternative; record what was and was not evaluated.
Incremental burden	Integration, compute, verification, human review, switching, reserves and validity horizon.
Durable asset	Exportable mission record, understood rights, inspected model/program, negative history and restore drill.
Next bottleneck	The specific missing evidence or capability whose improvement would change the decision.

25.1 A transparent scenario instead of a promise

Consider a **purely illustrative worksheet**, not a product result: 12,000 eligible annual cases, a lower-bound scenario spread of CA\$15 per case, and annual deductions $C=40,000$, $R=20,000$, $H=25,000$, $I=10,000$, $D=15,000$. The adjusted scenario is $12,000 \times 15 - 110,000 = \text{CA\$}70,000$ per year. At a CA\$7 spread it is **–CA\$26,000**. Initial asset wealth, unrelated gains and unpaid prospective capacity are not added.

The supplied calculator recomputes this arithmetic from explicit inputs; it neither estimates the spread nor authenticates its lower bound. A scenario should remain UNKNOWN until the required quantities are supplied. A financially positive estimate cannot compensate for a failed hard gate.

25.2 Keep the strongest component, not the preferred story

A frozen predictor can remain the best component while its planner, resource contract or observation process needs revision. Test one change at a time where possible. A better prediction metric does not establish better policy outcomes. The historical ORCHESTRA manuscripts retained in the product discuss this distinction; this guide does not rerun their industrial comparison or credit their scenario values to 20.0.1. [8]

A new mission may benefit more from better observations, longer planning, a simpler comparator or a clearer resource contract than from more training. Record which component changed. Never credit an institutional layer with the entire return of a predictor you could have used without it.

26 A portfolio that can grow without outrunning evidence

The most valuable use of this system is not to automate every possible action. It is to accumulate **mission-specific executable knowledge and a disciplined replacement process** around repeated consequential decisions. The product supplies bounded local mechanisms; a larger operating institution requires additional evidence, people and infrastructure. [2, pp. 43–46, 6, pp. 34–36]

Consider a fictional owner of a substantial industrial portfolio. The owner has capital, technical staff and many candidate missions, but no established return from this software. A rational programme does not multiply a flattering simulator score by all assets. It releases effort in stages:

Stage	Useful retained asset	Evidence required before further commitment
One operating question	Mission dossier, data map, serious alternatives	Observable outcomes and lawful evidence access
Local classifier or predictive model	Parameters/program, local comparison, failures	Independent evidence adequate to the real decision
Coordinated resource proposal	Exact predictor–planner composition	Complete policy outcomes, resource accounts and stress tests
Controlled operational pilot	Qualified interface, human decision and recovery	Separately accepted production assurance and permission
Reinvestment decision	Actual receipts, costs, reserves and delivery contract	Paid, delivered, usable capacity—not an NPV headline
New mission family	Retained useful structure and negative history	Fresh transfer evidence and renewed qualification

The research budget purchases decisions, not grades. A clear rejection may prevent further expenditure. Its avoided cost is not automatically a booked saving: the counterfactual must be credible and attributable before such a claim is made.

26.1 The strongest learning can be a change outside the model

When prediction improves but utility does not, test the planner, observation timing, terminal costs, evaluation design and human workflow. When a new model is tied, retain the simpler or better-supported component. When the best rental alternative wins, use that fact to change the formation decision rather than weakening the benchmark.

The larger horizon connects reliability, yield, logistics, materials, energy, computation and scientific experiments. Each link introduces constraints that cannot be paid away with rhetoric. A manufacturing plan needs feedstock and qualified equipment; a power system needs conversion and heat rejection; remote operations need transport, repair and communication delays. The following analytical chapters show how to make those dependencies explicit. They do not assert that this release implements that whole trajectory.

A **next-horizon record** should name the bottleneck, competing explanations, smallest distinguishing observation, full cost, owner, expiry, stop condition and evidence required for a subsequent mission. The institution advances when that record survives a serious comparison—not merely when another model is trained.



ALPHA FRONTIER

SUCCESSOR Ω

Ω
WORLDS
PEOPLE
IDEAS
INDUSTRY
A BRIGHTER
TOMORROW

Ω
EXPLORE
SETTLE
BUILD
PROTECT
BELONG

INTELLIGENCE
EXPANDS
POSSIBILITY

Qualify the next horizon

Conceptual editorial fiction. Not a product interface, measurement, proof or permission.

PART / PARTIE V

The advanced institutional companion

27 Advanced operator note: prediction, optimization and value

Conditional theory—not an additional product feature. A learned predictor estimates a transition $\hat{f}$. A planner chooses actions under that estimate. The evaluator observes outcomes from a separate source. These are different functions.

For a common feasible action set, suppose the true horizon value is $Q(a)$, the modeled value is $\hat{Q}(a)$, and $|Q(a) - \hat{Q}(a)| \leq \varepsilon$ uniformly. If the selected action has surrogate suboptimality at most g , then

$$Q(a^*) - Q(\hat{a}) \leq 2\varepsilon + g.$$

Derivation. Add and subtract $\hat{Q}(a^*)$ and $\hat{Q}(\hat{a})$. The two model-error terms are bounded by ε and the remaining optimization term by g . Smaller search error helps only within the surrogate; the bound requires a uniform model-value error not established by a small mean prediction error.

For a Lipschitz augmented-state transition with constant L , uniform one-step model error ϵ and a common start, $e_{k+1} \leq L e_k + \epsilon$ gives $e_k \leq \epsilon \sum_{j=0}^{k-1} L^j$. Delayed consequences can make a longer horizon useful while model error accumulates. Threshold penalties are discontinuous; the smooth bound alone does not cover them. These qualifications are central to model-predictive control, not optional disclaimers. [9]

Operator consequence: keep the model fixed when diagnosing a planner change, and keep the planner fixed when diagnosing a model change. Use fresh complete trajectories for an operating-value claim. Do not multiply correlated hourly records into a fictitious independent sample size.

28 Advanced operator note: scarcity and strategic behavior

Conditional theory. 20.0.1's process cells are centrally configured assets, not independent strategic owners. The following mathematics identifies what would need a separate contract or implementation.

28.1 Shadow prices explain a constrained surrogate

For action value $v_i(a)$, resource vector $r_i(a)$ and shared budget b , the discrete surrogate is

$$V_{\text{IP}} = \max_{a_i} \sum_i v_i(a_i) \quad \text{subject to} \quad \sum_i r_i(a_i) \leq b.$$

For every $\lambda \geq 0$, the Lagrangian expression

$$D(\lambda) = \lambda^\top b + \sum_i \max_a \{v_i(a) - \lambda^\top r_i(a)\}$$

upper-bounds a feasible allocation. A continuous relaxation can provide a dual certificate, but fractional actions are not factory instructions. A multiplier denominated in surrogate value per MWh is not a supplier's price, a site-day marginal saving or evidence of a market equilibrium. [10]

28.2 A truthful mechanism requires its assumptions

Under quasi-linear utilities and a fixed report-independent feasible set, exact welfare maximization plus the VCG externality payment makes truthful reporting a best response. Replacing exact welfare maximization with an arbitrary heuristic can invalidate the proof. 20.0.1's ratio score is not a VCG mechanism; there are no such participant payments in the product. [11]

28.3 Verification also has incentives

In an illustrative risk-neutral game, an unsupported claimant gains G if undetected, loses that gain and pays enforceable loss L if caught, and faces detection probability q . Deception has relative expected payoff $(1 - q)G - qL$. Truthful reporting is a best response only if $q \geq G/(G + L)$. A signature does not establish q , and a nominal penalty is ineffective when unenforceable.

Operator consequence: name the evaluator, disclose conflicts, preserve failed trials and avoid pass-contingent incentives. A local colleague or a second key is not automatically the independent organization required by the architecture. [2, pp. 17, 28–29]

29 Analytical laboratory: test the decision logic itself

Retained analytical companion. These conditional examples were introduced in Guide 4.0. Their assumptions, equations and numerical inputs are retained; this illustrated edition does not turn them into a new product-benefit experiment.

Status: new conditional teaching calculations, not product performance results. The examples below are executed by `08_TOOLS/analysis.py`. They use stated numbers, not customer observations. They add no planner, market mechanism, physical model or production authority to 20.0.1.

The purpose is to make five recurring decisions more precise: what to compare, what evidence to buy, which mechanism to trust, which physical bottleneck binds, and whether reinvestment is even feasible.

29.1 A declared metric changes the best direction

The Navigator and Gym use a mission-dependent architecture metric as a design analogy. For a **strictly positive-definite** metric G and local value gradient g , consider the budgeted directional problem

$$\max_{\delta} g^\top \delta \quad \text{subject to} \quad \delta^\top G \delta \leq b^2.$$

With $g \neq 0$, its solution is

$$\delta^* = \frac{bG^{-1}g}{\sqrt{g^\top G^{-1}g}}, \quad g^\top \delta^* = b\sqrt{g^\top G^{-1}g}.$$

To verify this, set $y = G^{1/2}\delta$ and apply Cauchy–Schwarz to $(G^{-1/2}g)^\top y$. Equality occurs when those vectors are parallel. In the teaching example $G = \text{diag}(9, 1)$, $g = (3, 2)$ and $b = 1$, the optimal step is approximately $(0.1491, 0.8944)$ and the directional gain is $\sqrt{5}$.

Operating consequence. A change with a larger technical score can be a worse use of scarce institutional effort. Improving evidence quality or recoverability may be preferable to changing the model. These coordinates and costs are assumed; the product does not estimate a corporate metric tensor.

The source’s more general positive-semidefinite metric need not be invertible. The derivation here explicitly adds strict positive definiteness. A pseudoinverse, regularization or a quotient-space construction would be a separate modeling choice, not a silent correction of the source or an implemented R9/R10 feature. [6, pp. 5–7, 7, pp. 14–16]

29.2 Allocate verification effort against a strategic choice

Suppose an attacker chooses between two pathways with declared unmitigated losses $L_1 = 100$ and $L_2 = 40$. Inspection effort gives detection probabilities q_1, q_2 at linear cost with $q_1 + q_2 \leq 1$. The owner minimizes

$$\min_{q \geq 0, q_1 + q_2 \leq 1} \max\{100(1 - q_1), 40(1 - q_2)\}.$$

At the interior optimum, use the full budget and equalize the two residual losses. Thus $q_1 = 5/7$, $q_2 = 2/7$, and worst residual loss is $200/7 \approx 28.57$. Spending everything on the higher-loss pathway leaves a loss of 40 on the other one.

This is an illustrative minimax calculation, not measured detection performance. Real detection probabilities can be unknown, nonlinear or correlated. The actor controlling inspection may have different incentives from the owner. In a simple reporting game, a false claim with private gain B is deterred only if its expected enforceable consequence qF offsets that gain, after all relevant participation and effort costs. A hash is not the penalty, an audit promise is not detection, and contractual language alone does not supply enforcement.

Operating consequence. Map distinct failure pathways and their evidence. Do not spend the whole review budget on what is easiest to measure while leaving an unexamined route to unsupported authority. The source’s producer–verifier–admitter separation motivates the question; this numerical game is a new teaching model. [6, pp. 9–11, 7, pp. 17–19]

29.3 A heuristic plus a familiar payment formula can fail truthfulness

The VCG truthfulness argument requires its allocation rule and payment construction to satisfy the stated assumptions. It does not survive arbitrary substitution of a heuristic. [11]

A capacity-three teaching problem has participants A, B and C. They require $(2, 3, 1)$ resource units and value service at $(4.1, 6, 0.1)$. A density-greedy allocator selects A then C: value 4.2, versus 6 for B alone. A pivot-style charge to A using the exact best welfare of the others is $6 - 0.1 = 5.9$. A’s true utility is therefore $4.1 - 5.9 = -1.8$.

By reporting zero instead, A is not selected and obtains zero utility. Truthful reporting is not a dominant strategy in this hybrid rule. This explicit counterexample does not mean all approximations are incompatible with incentives. It means that the allocation, payment rule, information assumptions and feasible set must be analyzed together.




Figure 5. A teaching distribution can lose its apparent margin as unfavorable retained cases receive more weight. This is not a stress test of product benefit.

Operating consequence. A centrally governed composition planner is not a truthful auction. Do not attach payment or autonomy rights to its scores without constituting and testing the additional mechanism. No such payments are executed by this guide.

29.4 Stress the distribution, not only the model

Consider a fixed, explicitly hypothetical vector of sixteen paired advantages,

$$d = (-8, -4, -2, 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12),$$

with uniform reference weights $p_i = 1/16$ and a declared deduction of two utility units. The nominal gross mean is four. What happens when an adverse mixture of these *same* outcomes becomes more common?

Define the support-limited robust value

$$V(\rho) = \min_{q \in \Delta: D_{\text{KL}}(q \| p) \leq \rho} \sum_i q_i d_i.$$

For an interior radius, the minimizing weights are $q_i \propto p_i e^{-d_i/\eta}$ and the dual expression is

$$V(\rho) = \sup_{\eta > 0} \left[-\eta \log \sum_i p_i e^{-d_i/\eta} - \eta \rho \right].$$

The script solves the multiplier by bisection and checks primal–dual agreement over 101 radii. The witness changes probabilities, not outcome records. It is **not a confidence interval**, an estimated future distribution or protection against failures outside the retained support. [10]

Operating consequence. Before scaling a favorable estimate, ask how much distributional change the decision can tolerate, which failure types were never represented, and whether a reserve merely disguises a failed hard gate.

30 Advanced operator note: physical accounts before large-scale promises

Conditional theory. The composition adapter's normalized heat state is not itself temperature or entropy. Even a simulator with a closed numerical energy balance is not a calibrated industrial thermodynamic model. A physical mission requires measured units, material properties, boundary conditions and conservation tests.

For a closed system with stored energy U , entropy S , electrical input P_e , useful output P_u , heat ports $\dot{Q}_k$ at temperatures T_k and reference T_0 ,

$$\dot{U} = P_e - P_u + \sum_k \dot{Q}_k, \quad \dot{S} = \sum_k \frac{\dot{Q}_k}{T_k} + \dot{S}_{\text{gen}}, \quad \dot{S}_{\text{gen}} \geq 0,$$

so

$$\frac{d}{dt}(U - T_0 S) = P_e - P_u + \sum_k \left(1 - \frac{T_0}{T_k}\right) \dot{Q}_k - T_0 \dot{S}_{\text{gen}}.$$

Open factories require material-flow exergy terms. Intelligence may improve allocation or reduce avoidable losses; it cannot remove the last term by improving a reward. A passive thermal network with heat capacities $C_i > 0$ and symmetric conductances $\kappa_{ij} \geq 0$ obeys

$$C_i \dot{T}_i = \sum_j \kappa_{ij}(T_j - T_i), \quad \frac{d}{dt} \sum_i C_i T_i = 0, \quad \dot{S} = \sum_{i < j} \kappa_{ij} \frac{(T_i - T_j)^2}{T_i T_j} \geq 0.$$

Pairwise flux cancellation proves energy conservation; pairing the entropy terms proves nonnegative entropy production. This is a useful structural check for a *future physical adapter*, not a claim that all learned 20.0.1 programs satisfy it. [12]

For useful electrical power P , collection efficiency η , incident irradiance $I_\star$ and radiating temperature T_r ,

$$A_c \geq \frac{P}{\eta I_\star}, \quad A_r \geq \frac{P_{\text{heat}}}{\epsilon \sigma (T_r^4 - T_{\text{bg}}^4)}.$$

If all absorbed energy ultimately thermalizes inside the chosen boundary, $P_{\text{heat}} = P/\eta$, not merely P . Conversion losses must be included once in the correct boundary. At a deliberately idealized $I_\star = 1361 \text{ W/m}^2$, $\eta = 0.30$, $\epsilon = 0.90$, $T_r = 340 \text{ K}$ and $T_{\text{bg}} = 3 \text{ K}$, one useful TW requires approximately $2,449 \text{ km}^2$ of collector area and $4,888 \text{ km}^2$ of total-system radiating surface. These are reference identities, not a design or a construction forecast. [13, 14]

The future mission must also price material mass, qualified manufacturing, heat rejection, transport, repair and communications latency. A high-value software result supplies none of those quantities automatically.

31 Physical integrity: a small exact system before an enormous horizon

Status: physical theory and new numerical illustrations, not learned industrial physics. A normalized state variable inside a workbench is not automatically kelvin, joules or conserved inventory. Domain experts must establish that mapping separately.

31.1 Preserve structure before optimizing a reward

An isolated thermodynamic model can be written schematically as

$$\dot{z} = L(z)\nabla E(z) + M(z)\nabla S(z),$$

with $L^\top = -L$, $M \geq 0$, $L\nabla S = 0$ and $M\nabla E = 0$. These conditions imply

$$\dot{E} = 0, \quad \dot{S} = \nabla S^\top M \nabla S \geq 0.$$

The equalities follow from antisymmetry and the degeneracy conditions. Without those conditions, writing “energy” and “entropy” on two variables proves nothing. For open systems, power, heat, material and entropy fluxes must be included explicitly. The supplied product does not enforce this formalism merely because its planner has a resource limit. [12]

31.2 Two thermal bodies: an exact reference

Let two bodies with constant positive heat capacities C_1, C_2 exchange heat through conductance $k > 0$, with no external flow:

$$C_1 \dot{T}_1 = -k(T_1 - T_2), \quad C_2 \dot{T}_2 = k(T_1 - T_2).$$

Then $E = C_1 T_1 + C_2 T_2$ is constant and

$$\dot{S} = k(T_1 - T_2) \left(\frac{1}{T_2} - \frac{1}{T_1} \right) = \frac{k(T_1 - T_2)^2}{T_1 T_2} \geq 0.$$

For $C_1 = 2$ MJ/K, $C_2 = 3$ MJ/K, initial temperatures 500 K and 300 K, and $k = 1,200$ W/K, equilibrium is **380 K**, the decay time is **1,000 seconds**, and total modeled energy is **1.9 GJ**. The script evaluates the closed-form trajectory at 241 times and checks energy conservation and monotone entropy.

Operating consequence. Use such independently specified reference cases to challenge a future physical adapter. Do not let a learned transition redefine the balance by which its feasibility is judged.




Figure 6. The exact passive thermal reference approaches a common temperature without creating energy. It is an analytical illustration, not an R10 learned world.

31.3 Information has a physical account—not a supernatural multiplier

For a classical finite-state system at temperature T , with canonical equilibrium $\pi_i \propto e^{-E_i/(k_B T)}$, define

$$F[p] = \sum_i p_i E_i + k_B T \sum_i p_i \log p_i.$$

Direct substitution gives $F[p] - F[\pi] = k_B T D_{\text{KL}}(p||\pi) \geq 0$. This concerns a physical distribution and its equilibrium reference, not an arbitrary business score. Measurement and feedback can change work bounds only with a complete information-processing account. [15]

Erasing one unbiased classical bit in the ideal quasistatic setting has the familiar lower heat bound $k_B T \ln 2$. At 300 K it is approximately 2.871×10^{-21} J. This does not predict the energy consumed by a model, database, cryptographic signature or computer. Actual devices have many other losses. No computational efficiency claim is inferred from the bound. [16, 17]

31.4 Count conversion heat as well as useful-power dissipation

For a declared electrical payload P supplied at conversion efficiency η , suppose the delivered electrical power is eventually dissipated locally. The payload heat is P , conversion heat is $(1/\eta - 1)P$, and total rejected heat is P/η .

Under an idealized isothermal gray-surface model with effective radiating area A , emissivity ϵ and view factor one,

$$A \geq \frac{P/\eta}{\epsilon \sigma (T_r^4 - T_b^4)}.$$

This expression excludes transport losses, nonuniform temperature, shading, pumping, protection, structural support and degradation. The new table evaluates 1 MW, 1 GW and 1 TW at four radiator temperatures, with explicitly assumed efficiency and areal mass. These are **resource requirements under assumptions**, not feasible designs. [14, 18]




Figure 7. Heat-rejection area grows with the assumed payload, but that algebra does not qualify fabrication or deployment at the new scale.

Operating consequence. When the ambition expands, the next mission is to measure the binding constraint: specific mass, lifetime, conversion yield, material supply, manufacturing rate, thermal transport or maintenance. Funding cannot replace any missing physical capability.

32 Advanced operator note: productive expansion needs a network

Conditional theory. Let $K(t)$ denote a vector of productive capacity types, D a positive diagonal replacement-rate matrix, $B \geq 0$ an irreducible matrix of fully paid qualified additions, and τ their common delivery delay:

$$\dot{K}(t) = -DK(t) + BK(t - \tau).$$

A positive exponential mode exists precisely when $\rho(D^{-1}B) > 1$. Substitution of $K(t) = e^{rt}v$ gives

$$\rho((rI + D)^{-1}B)e^{-r\tau} = 1.$$

For $r \geq 0$ the left-hand side decreases continuously; it crosses one at a positive rate only under that threshold. Increasing a finite delay lowers the positive rate. A hard delivery ceiling $\bar{u}_i$ instead implies

$$K_i(t) \leq e^{-\delta_i t} K_i(0) + \frac{\bar{u}_i}{\delta_i} (1 - e^{-\delta_i t}).$$

Thus capital alone does not abolish a manufacturing bottleneck. The model is an idealization; no coefficient of B is measured by the tutorial. It is not a projection of a customer's growth.

Practical translation: pursue a sequence of different qualified missions. First improve a repeated operating decision. Then investigate the limiting material, energy or delivery mechanism. Only realized cash can pay for capacity. Preserve the methods and failures so the next team need not confuse old evidence with a newly established result.




Figure 8. Longer commissioning delays reduce the conditional expansion factor. Neither the operator nor its parameters are estimated by the product.

33 Productive expansion: the threshold is conditional, delivery is real

A long-lived intelligence institution is useful when it helps choose and requalify successive productive missions. It does not make a claim of self-financing growth true merely by retaining a Chronicle.

Let $K_t \in \mathbb{R}_+^n$ represent separately measured productive capacities. Under an explicitly idealized, homogeneous model,

$$K_{t+1} = (1 - \delta)K_t + BK_{t-d},$$

where $B \geq 0$ converts earlier productive activity into **delivered** new capacity, d is commissioning delay and $\delta > 0$ is proportional loss. For a positive irreducible B , consider its Perron direction $Bv = \rho(B)v$. A solution $K_t = r^t v$ satisfies

$$r^d(r - 1 + \delta) = \rho(B).$$

The positive root exceeds one exactly when $\rho(B) > \delta$. For a given excess above replacement, a longer delay reduces the growth factor. This statement concerns the specified recurrence; finance, prices, resources, demand and rights do not appear by magic inside B .

The teaching matrix

$$B = \begin{bmatrix} 0.08 & 0.03 \\ 0.04 & 0.07 \end{bmatrix}, \quad \delta = 0.05$$

has $\rho(B) = 0.11$. The script solves commissioning delays from zero through eight steps and verifies the characteristic equation. If instead the delivered increment is bounded by $\bar{D}$, a componentwise inequality $K_{t+1} \leq (1 - \delta)K_t + \bar{D}$ bounds each corresponding capacity by its initial stock and $\bar{D}/\delta$. A fixed manufacturing ceiling defeats indefinite geometric expansion.

33.1 Keep cash distinct from a productive-capacity operator

A separate illustrative account assumes an initial incremental gross annual effect of CA\$12 million, erosion of 8% per year, CA\$3 million in deductions, a 50% distribution policy, CA\$10 million per module and a two-year lead time. It retires the incremental policy when it no longer covers deductions. New modules receive **no invented additional earnings**.

The script checks, year by year,

$$\text{opening cash} + \text{available surplus} = \text{distributions} + \text{paid capacity} + \text{closing cash}.$$

All inputs are teaching assumptions. The table is neither a forecast nor a valuation. Tax, financing, demand, operating performance and delivery evidence require their own work. The useful point is the separation among valued opportunity, realized receipts, reserves, payment and commissioned capacity.

A productive research ladder: operating reliability → material yield → fabrication throughput → energy and computation → remote production. Each arrow means another mission, interface, evidence requirement and stop condition. The guide provides a method for pursuing that horizon without claiming that the present product has already reached it.

PART / PARTIE VI

Reference and evidence

34 Command reference: copy a complete command, then inspect its result

The authoritative commands are in the actual 20.0.1 executable. The current 07_NATIVE/README_EN_FR.md retains an R9 heading; use **current CLI help**, VERSION.json and the version-20 operating guide to resolve that historical label. The old heading does not indicate that the extracted product is R9.

03_COMMANDS/COMMANDS_EN.txt contains one command per line; the French catalogue preserves the same machine syntax. PDF blocks use POSIX \ continuations where needed. Never break a filename token to fit a line. Replace uppercase placeholders with values actually returned by the product.

34.1 Orientation

```
python3 VERIFY_PACKAGE.py
node 07_NATIVE/start.mjs doctor
node 07_NATIVE/start.mjs help
node 07_NATIVE/start.mjs mission help
node 07_NATIVE/start.mjs compose help
node 07_NATIVE/start.mjs mission wizard
```

There is **no world help command in this release**. The general help enumerates the World Engine commands. A refused unknown command is not a failed model.

34.2 Prepare without training

```
node 07_NATIVE/examples/v20/prepare.mjs \
  --route compose --out ../manual-mission \
  --language en --synthetic
```

Choose classify or world instead for those formats. This operation prepares public synthetic inputs and a HOLD draft only.

34.3 Confirm and form

After completing and binding the dossier, the composition route is:

```
node 07_NATIVE/start.mjs compose preflight \
  --plan ../manual-mission/BOUND.json
node 07_NATIVE/start.mjs compose confirm \
  --plan ../manual-mission/BOUND.json --token TOKEN_FROM_PREFLIGHT \
  --out ../manual-mission/CONFIRMED.json
node 07_NATIVE/start.mjs compose form \
  --plan ../manual-mission/CONFIRMED.json --out ../manual-run
```

Without a secret flag, the native formation command requests hidden input. Use `--passphrase-stdin` or `--secret-fd` only through an appropriately protected process route; do not put the passphrase in command arguments. The guide's teaching companion uses a private pipe. Parent-secret descriptors are separate when authorizing a continuation.

For a classifier, omit `compose` and use `create`. For a temporal program, use `world preflight`, `world confirm` and `world synthesize`. For a child, the classifier command is `successor --file PARENT --plan ...`; the world command is `world revise --parent PARENT --plan ....` Composition formation reads the parent from its constituted plan.

34.4 Verify, pin and propose

```
node 07_NATIVE/start.mjs compose verify \
  --file ../manual-run/composition.successor.json
node 07_NATIVE/start.mjs compose pin \
  --file ../manual-run/composition.successor.json \
  --trust ../operator-trust.json --expected EXACT_PAYLOAD_DIGEST \
  --actor "Local research owner" --reason "Reviewed synthetic scope"
node 07_NATIVE/start.mjs compose plan \
  --file ../manual-run/composition.successor.json \
  --trust ../operator-trust.json --context ../CONTEXT.json \
  --out ../NEW_PROPOSAL.json
```

A **context must actually exist**, match the adapter, be within scope and remain timely. The v20 preparer does not silently fetch live observations. `08_TOOLS/make_synthetic_context.mjs` prepares a fresh public context for the default synthetic fleet only; it is not an arbitrary-customer adapter and must never retimestamp real evidence.

For worlds, the syntax differs: `world predict --candidate FILE --context FILE --actions 0,1` and `world plan --candidate FILE --context FILE`. Do not substitute the composition's `--file` flag into those commands.

34.5 Return to an existing result

Use the reader appropriate to the artifact. The general `verify --file` dispatches current artifact types; family-specific readers expose their own result shape. For a composition, `compose status --file FILE --trust TRUST.json` adds the current operator journal and clock to the historical record. No background machine monitor is started.

35 Evidence, provenance and what was tested for this guide

The exact unchanged customer archive is `SUCCESSOR_OMEGA_20_0_1_CUSTOMER_PRODUCT.zip`: **52,044,463 bytes**, SHA-256 `bc9ff32b3b9640183251fad2c84e42782ae13b164ea0bfddc1d1fe0dbeaf6d6e`. Guide source, teaching tools and analytical examples have separate identities. Edition 4.1 preserves the product ZIP and the Guide 4.0 numerical and teaching tools unchanged. No new product release, contract or earned designation is created here.

35.1 Retained Guide 4.0 execution record

Source-reported Guide 4.0 check	Scope and result
Exact unchanged native product	171 native tests; 33 core self-test checks passed
Integrated mission terminal	46 English/French checks passed
New teaching companion	Six actual language/route journeys; 45 checks passed

Source-reported Guide 4.0 check	Scope and result
Additional native walkthrough	Ten CLI calls; seven outcome/continuity checks passed
Worked refusal	Explicit reuse reports insufficient eligible evaluation evidence
Conditional analytical laboratory	19 test groups; no observed customer benefit

These are the source-reported Guide 4.0 results. The groups overlap; their counts are not independent observations of reliability or benefit. Historical evidence is retained in 05_EVIDENCE_HISTORY_4_0/. The new illustrated-edition checks, source build and final archive verification appear in 05_EVIDENCE/; old results are not relabeled as newly executed.

The historical Guide 4.0 machine-readable acceptance record identifies its actual product-inventory checks, retained native regression suite, core self-test, integrated native terminal journeys, new teaching-companion journeys, synthetic refusals, arithmetic checks, and document builds. Counts are reported only after execution. **The installed product was not patched to make the guide pass.**

The retained lesson evidence is public synthetic material. The companion performs actual preflight, confirmation, formation, signing, verification, export and restore calls. A successful private-pipe test does not establish a hostile-host secret boundary. A same-host restoration does not establish external organizational independence.

The delivered historical product acceptance is preserved separately under 07_REFERENCES/. Its release tests are not relabeled as new guide tests. A paper may accurately describe a source-reported result without having rerun it; the register makes that distinction explicit.

35.2 Scope of this illustrated publication

Neither the retained campaign nor this illustrated revision establishes a human novice-usability study, macOS or Windows support, current-patched-runtime qualification, real Docker operation, live merchant checkout, external independent security review or measured industrial benefit. Native browser navigation is assessed only where actually attempted; in-memory controls do not substitute for it.

The existing sale-readiness questions are not closed by a new guide. In particular, a publication does not repair a merchant's version-binding terms, qualify an untested supported installation or exercise a live checkout. This package grants no new platform licence and supplies no commercial warranty beyond the original agreement.

35.3 How to reproduce the publication

02_SOURCE/build.py generates the English/French A4 and Letter PDFs plus first-session cards from editable Markdown and the supplied LaTeX template. The generated .tex files are included. 08_TOOLS/analysis.py reproduces the explicitly conditional teaching calculations. VERIFY_GUIDE.py verifies the final inventory. Fonts are referenced by standard installed names; font files are not redistributed.

Full reading editions provide native HTML headings and MathML equations, offline links and one-line command copying. No PDF/UA or assistive-technology certification is asserted. Concept art is separated from programmatic diagrams and is non-authoritative.

36 Glossary

Term	Practical meaning
A0	Local observation, research and informational proposals; no production authority granted here.
Adapter	A precise mapping between data and the numerical model the product supports.
Alpha	A scoped advantage after the specified uncertainty and incremental deductions; local arithmetic is not independent proof.
Beta	The constituted alternative; a local proxy is not an externally evaluated frontier model.
Calibration	A specified procedure estimating predictive uncertainty under assumptions.
Candidate	A particular fitted model or executable hypothesis.
Composition	The complete predictor, planner, objective, resources, calibration, scope and evidence identity.
Confirmation	Explicit acceptance of a resolved plan and its input commitments before execution.
Chronicle	Ordered, tamper-evident local history; not truth or independent proof by itself.
Exposure	Known prior use or inspection of evidence that limits freshness claims.
Freeze	Binding the exact composition before the corresponding evaluation.
I0	Internal/local evaluation within the claimant's workflow.
Pin	An explicit local operator acceptance of an exact identity.
Proof request	A specification asking for evaluation; not a proof receipt.
Rehearsal	Explicit synthetic practice, ineligible for automatic preferred promotion.
Restore	Recreate retained artifacts and history under the rights boundary; not inherited trust.
Scope	The population, time, resource and uncertainty conditions under which a local proposal is permitted.
Successor	A descendant whose changes, evidence and local decision are recorded; improvement is not presumed.

Retain the mission. Inspect the evidence. Improve one decision at a time.

SUCCESSOR Ω / USER GUIDE Ω

Knowledge before expansion



Conceptual editorial fiction. Not a product interface, measurement, proof or permission.

References

- [1] MONTREAL.AI / QUEBEC.AI. *SUCCESSOR Ω 20.0.1 R10.1: Customer Product*. Exact supplied customer ZIP: 52,044,463 bytes; SHA-256 bc9ff32b3b9640183251fad2c84e42782ae13b164ea0bfddc1d1fe0dbef6d6e. Product unchanged by this guide. 2026.
- [2] V. Boucher. *Neural-Symbolic SUCCESSOR Ω : Customer-Owned Mission Intelligence*. Definitive Corporate White Paper. 59 pages. Exact supplied architectural reference, SHA-256 36e7c26fe28961dbd8a0dc89160a0462e08acb53fc9b00e2508a8740b52847e9. Proposed architecture is distinct from release qualification. Aug. 25, 2026.
- [3] MONTREAL.AI / QUEBEC.AI. *20.0.1 native source, data contracts, composition guides and support terms*. Paths and hashes in 05_EVIDENCE/SOURCE_MAP.json. In particular 07_NATIVE/start.mjs, lifecycle, world, composition; 08_GUIDES and 04_LICENSE_AND_POLICIES_EN_FR. 2026.
- [4] Node.js Project. *Node.js downloads and release status*. A current release list is not an executed product compatibility test. URL: <https://nodejs.org/en/download> (visited on 09/06/2026).
- [5] MONTREAL.AI / QUEBEC.AI. *SUCCESSOR Ω 20.0.1 R10.1: Quick Start / Démarrage rapide*. Fourteen-page supplied bilingual quick start; current source controls exact commands. 2026.
- [6] V. Boucher. *Mission-Sovereign Successors: GoalOS Singularity Navigator Ω + SEIZE—Alpha, Beta, Spread, and the Verified Succession Institution in Singularity Time*. Version 5.0.0. Exact attached A4 edition: 71 physical PDF pages, 61 numbered main-body pages. Printed pages are used in citations. SHA-256 in the integration input register. Institutional design, not current product execution proof. Aug. 9, 2026.
- [7] V. Boucher. *Gym, Specialist ASI and Successor Ω : The Sovereign Mission Environment, the Proven Intelligence Institution, and the Successor to the Successor in Singularity Time*. Version 2.1.0. Exact attached US Letter edition: 83 physical PDF pages, 71 numbered main-body pages. Printed pages are used in citations. Synthetic Aurelia case is unaudited; source APIs are architectural pseudocode. Aug. 10, 2026.
- [8] V. Boucher. *SUCCESSOR Ω —ORCHESTRA Ω : Intelligence Under Shared Constraints*. Author edition 2.0, historical R8 study retained in the product publication bonus. Not a new product-20.0.1 economic evaluation. 2026.
- [9] J. B. Rawlings, D. Q. Mayne, and M. M. Diehl. *Model Predictive Control: Theory, Computation, and Design*. 2nd ed. Author-hosted second-edition resource page, sixth printing listed at consultation. General control background; no imported performance claim. Nob Hill Publishing, 2026. URL: <https://sites.engineering.ucsb.edu/~jbrow/mpc/> (visited on 09/06/2026).
- [10] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004. URL: <https://web.stanford.edu/~boyd/cvxbook/> (visited on 09/06/2026).
- [11] T. Roughgarden. *Incentives in Computer Science, Lecture 15: The VCG Mechanism*. Stanford University. Nov. 14, 2016. URL: <https://theory.stanford.edu/~tim/f16/l/l15.pdf> (visited on 09/06/2026).
- [12] A. M. Badlyan, B. Maschke, C. Beattie, and V. Mehrmann. *Open Physical Systems: From GENERIC to Port-Hamiltonian Systems*. Version 4. Background for structured energy, dissipation and exergy accounts; not a product experiment. 2018. arXiv: 1804.04064.
- [13] NASA Goddard Space Flight Center. *Solar Irradiance Science*. The 1361 W/m² reference is normal irradiance near Earth’s orbital distance, not globally averaged ground input. URL: <https://earth.gsfc.nasa.gov/climate/projects/solar-irradiance/science> (visited on 09/06/2026).
- [14] National Institute of Standards and Technology. *CODATA Fundamental Physical Constants*. Stefan–Boltzmann constant derived from SI defining constants; numerical value used: 5.670374419 times 10 to the minus 8 W m⁻² K⁻⁴. URL: <https://physics.nist.gov/cuu/Constants/> (visited on 09/06/2026).
- [15] T. Sagawa and M. Ueda. “Second Law of Thermodynamics with Discrete Quantum Feedback Control”. In: *Physical Review Letters* 100 (2008). Feedback-work background. The guide separately derives a classical finite-state free-energy identity., p. 080403. doi: 10.1103/PhysRevLett.100.080403.
- [16] A. Bérut, A. Arakelyan, A. Petrosyan, S. Ciliberto, R. Dillenschneider, and E. Lutz. “Experimental verification of Landauer’s principle linking information and thermodynamics”. In: *Nature* 483 (2012), pp. 187–189. doi: 10.1038/nature10872.
- [17] NIST. *CODATA: Boltzmann constant*. The SI value 1.380649 times 10 to the minus 23 J/K is exact. URL: <https://physics.nist.gov/cuu/Constants/> (visited on 09/06/2026).
- [18] NASA. *Small Spacecraft State of the Art: Thermal Control*. Physical control-design context; not a qualification of the guide’s illustrative radiator design. URL: <https://www.nasa.gov/smallsat-institute/sst-soa/thermal-control/> (visited on 09/06/2026).