

SUCCESSOR Ω

USER GUIDE Ω

From your first mission
to enduring intelligence.

SEIZE decisions · Mission Gyms · Tested local practice



CONCEPTUAL HORIZON / GENERATED ART / NOT A FACILITY OR ENGINEERING RESULT

Vincent Boucher

President, QUEBEC.AI & MONTREAL.AI

Constitute. Learn. Compare. Retain. Requalify.

Companion to unchanged SUCCESSOR Ω 19.0.0 R9. Local A0 research; no independent proof, production admission or realized Alpha.

Contents

1	Start here: one useful decision, one inspectable result	1
2	What you own, what you run, and what stays human	2
3	Choose the right workbench	3
4	Install, extract and orient yourself	4
5	Your first successful session	5
6	Configure a real mission before supplying real data	7
7	Workbench A: customer-specific classification	8
8	Read comparisons without fooling yourself	9
9	Create a successor deliberately	11
10	Workbench B: executable predictive worlds	12
11	Interpret uncertainty, contradictions and revision	13
12	Workbench C: compose intelligence under shared constraints	13
13	Pin, inspect and propose	15
14	Compare compositions and decide locally	16
15	Use SUCCESSOR Ω for the highest useful value	16
16	A practical operating cadence	17
17	From a local specialist to an admitted institution	18
18	Restore, move and protect your work	18
19	Troubleshooting without weakening the boundary	19
20	Academy, support and frequently confused features	20
21	Advanced operator note: prediction, optimization and value	21
22	Advanced operator note: scarcity and strategic behavior	22
23	Advanced operator note: physical accounts before large-scale promises	23
24	Advanced operator note: productive expansion needs a network	24
25	Command reference and tested copies	24
26	Source, verification and distribution notes	26
27	Glossary	28
28	Before formation: use SEIZE to choose the next decision	29
29	Constitute the Mission Gym without confusing it with a world model	31
30	Turn the mission into twenty-one bounded work-and-proof jobs	32
31	Preserve the distinct protocols, symbols and decisions	34
32	A worked mission: decide first, then use the supported route	35
33	Preserve failure, substitute carefully, and re-underwrite	37
34	Integration register: what changed and where it came from	38
	References	39

1 Start here: one useful decision, one inspectable result

Your first goal is not to build a universal autonomous organization. It is to complete one small, understandable learning cycle and retain the exact result. Start with the supplied synthetic evidence. Learn what each output means. Only then prepare evidence for your actual mission.

SUCCESSOR Ω 19.0.0 R9 is a local workbench with three distinct routes: a bounded classifier, an executable World Engine, and the Composition & Applicability Workbench. The retained Academy explains the institutional method. The white paper describes a larger target architecture; this guide does not substitute that architecture for the software that is actually delivered. [1–3]

Your immediate question	Read next	Useful first outcome
"I have never used this product."	Sections 2–5	Open the correct console and complete a synthetic rehearsal.
"I have business cases and preferred decisions."	Sections 6–9	A trained classifier with explicit comparison records.
"Actions have delayed consequences."	Sections 10–11	An executable temporal hypothesis and a hypothetical plan.
"Several assets share resources."	Sections 12–14	A frozen composition, an explicit trust pin and a local proposal.
"I need to judge value, not learn commands."	Sections 15–17	A mission brief, credible baseline and costed next decision.
"I need to recover or diagnose a refusal."	Sections 18–20	A checked restoration or an actionable error response.
"I need the mathematics."	Sections 21–24	The assumptions behind prediction, allocation, incentives and physical scale.

Three meanings of success. A command can finish successfully while its candidate fails comparison. A file can verify while its claim is out of date. A local result can be promising while independent proof and production admission remain absent. Read these fields separately; do not turn one green indicator into a universal verdict.

The quickest practical route. Read the first-session cards, extract R9 once, then launch `python3 START_PRACTICE.py` from this guide folder. The optional guide companion prepares explicit synthetic inputs and calls the unchanged product. It never trains, confirms a plan or pins trust without your separate choice. You can always use the native wizard directly.

New: choose the mission before increasing training. The attached GoalOS Navigator and Mission Gym papers are integrated in Sections 28–34, with a SEIZE decision card, an environment specification, the 21-job checklist and a source-to-R9 boundary map. A beginner can still start with the synthetic first session. An owner choosing a real mission should first read SEIZE and complete 09_MISSION_KIT/WORKSHEETS_EN.md. These source-derived planning tools do not grant authority or change the product. [4, 5]

Deliverable map. 09_MISSION_KIT/ contains the new planning records and 05_REFERENCE/ the exact attached papers. 04_EVIDENCE_GUIDE_2_0/ preserves earlier guide qualification; 04_EVIDENCE/ records this integration. 01_READ/ contains this guide and its reading editions. 03_EXERCISES/ contains worksheets and example records. 08_TOOLS/ prepares fresh synthetic exercises and verifies the guide. 06_PRODUCT/ contains the exact, unchanged R9 customer ZIP. The original product's licensing and customer rights remain unchanged.

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. [6, 3, pp. 36–37]

2.1 Three evidence badges

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 R9 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 R9 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. R9 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. [3, pp. 11–12, 33, 51–54]

3 Choose the right workbench

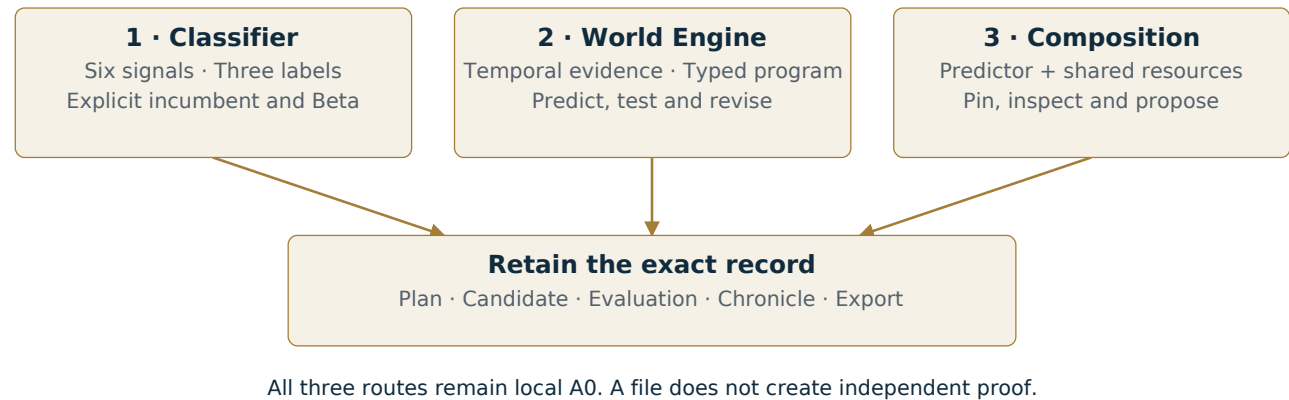


Figure 1. Choose by evidence representation, not by ambition. These are product routes, not proof levels.

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

Route	Required representation	What R9 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. [6]

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, extract and orient yourself

4.1 Preserve the original download

Keep the ZIP unchanged as your reference copy. Extract into a short path you control, such as a folder named `Omega` under your home directory. Do not run from an archive preview, an email attachment preview, a cloud placeholder or a read-only directory. Keep your research outputs in a separate working directory.

In this guide package, open `06_PRODUCT/` and extract `SUCCESSOR_OMEGA_19_0_0_CUSTOMER_PRODUCT.zip`. The folder you need contains `VERSION.json`, `VERIFY_PACKAGE.py`, `START_HERE.html` and `07_NATIVE/`. This is the **customer root**. The guide's `START_HERE.html` is a reading console; the product's `START_HERE.html` is its entry page.

4.2 A folder layout you can recognize

Use separate sibling folders. The guide and product may both contain a file named `START_HERE.html`; only the product contains `07_NATIVE/`.

Omega/	
USER_GUIDE_2_0/	reading, source and practice companion
omega-product/	
SUCCESSOR_OMEGA/	original R9; do not write results here
omega-practice/	synthetic inputs and your teaching runs

From the **guide** folder, this helper verifies the embedded ZIP before extracting it. The destination must not already exist:

```
python3 08_TOOLS/unpack_product.py --out ../omega-product
```

To open a terminal in a folder: use the folder's terminal action when your desktop provides one, or type `cd`, a space, followed by the quoted absolute folder path. In Windows PowerShell, for example, `cd "C:\Users\YourName\Omega\USER_GUIDE_2_0"` selects the guide. That path is an illustration: substitute your actual location. On macOS/Linux the path usually starts with `/`, not a drive letter. The displayed shell prompt is not part of a command.

Copy reliability. Inline command spaces are preserved in this edition. Long printed commands can still wrap across lines: use `01_READ/COMMANDS_EN.txt` or the HTML reading edition for a literal single-line copy. Do not paste typographic line-continuation arrows, a heading or a shell prompt. Host instructions are not Windows/macOS execution qualification.

4.3 Check the runtime

Node.js runs the product. Python 3 verifies packages. The optional guide helpers require Python 3.11 or later. Core use needs no npm installation, model subscription, API key, wallet or GPU. An internet connection is not required after the software and needed runtimes have been installed. [6]

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

On Windows, `py -3` may be the installed Python launcher [7]; substitute it for `python3` when appropriate. Open a terminal in the customer root. A quick check is that `07_NATIVE/start.mjs` exists there. Quotes protect paths containing spaces.

Use an organization-approved, supported and security-patched LTS runtime. The Node release pages consulted for this guide identify current LTS lines; the recorded R9 test patch is **not** an installation recommendation. Installing a current runtime does not by itself qualify R9 on that runtime. Run the relevant checks and retain the exact version. [8, 9]

Host evidence: this guide's executions use Linux x86_64, Node.js 22.16.0 and Python 3.13.5. Windows, macOS and a current patched LTS are not newly qualified here. Those limits do not prevent reading the PDFs or HTML on other devices. The iPhone/iPad reading experience is not a native execution qualification.

4.4 Verify before researching

From the customer root:

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

`VERIFY_PACKAGE.py` checks the delivered inventory. `doctor` identifies the local runtime and boundary. `selftest` is the retained core/Academy check, not every newer composition test. The broader local regression command is in the technical reference. A hash confirms identity relative to a known digest; it is not an independent security audit.

The exact customer ZIP used by this guide has **43,940,303 bytes** and SHA-256:

```
0fa222c25ba7a52e2e28c05aaa706c11ed01cdbacdfdeaff263ab15fb5632872
```

4.5 Open the console

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

Choose `en` or `fr`. Menu **1** creates a classifier mission; **2** creates its successor; **3** verifies; **4–6** manage the classifier's local decision, export and restore; **7** opens the World Engine; **8** opens the retained Academy; **9** opens compositions. Use each route's own commands for its decisions and restoration. The same-looking file extension does not make every route interchangeable.

5 Your first successful session

5.1 The guided practice route

Start with no confidential evidence. The guide companion is an optional teaching interface around the original product—not a replacement R9 application. It supports a new synthetic session and a later return to that session. It does not resume an interrupted optimizer invisibly.

From the **guide** folder:

```
python3 START_PRACTICE.py
```

Choose en or fr, supply the extracted R9 folder containing 07_NATIVE/, and choose a new practice folder outside the product. Full paths work, including spaces and accented characters. The companion checks the product files against the exact embedded archive before calling its code.

Choose **3: composition** for the first complete lesson. The other choices are **1: classifier** and **2: World Engine**. These are companion menu numbers, not the native product's menu numbers.

Stage	Your deliberate action	What the software actually does
Prepare	Choose a new practice directory.	Writes explicit synthetic data and plans; no training.
Inspect	Read the native preflight: evidence, roles, scope, predictor, budgets, unknown costs and A0.	Calls R9's non-training preflight.
Confirm	Type CONFIRM only for the displayed plan.	Calls the native confirmation with that exact token.
Form	Type TRAIN, then enter and repeat a private passphrase.	Passes the secret through standard input, not command arguments or a saved file; runs native bounded formation.
Verify	Read the original verification and result.json.	Reopens the signed file using R9. Integrity is not mission success.
Pin	For composition, explicitly type PIN after inspecting the full identity; name the operator.	Calls the product's separate local trust operation.
Propose	Type PROPOSE to generate a new explicit synthetic context.	Calls the original example generator and pinned planner; no real observation is retimestamped.
Restore	Type RESTORE.	Calls native export, restore and verification; checks that signed bytes match. No restored trust or preference is assigned.

The French confirmation words are CONFIRMER, FORMER, ANCRER, PROPOSER and RESTAURER. The optional world-plan confirmation is PLAN in either language. Type :cancel at a companion question to return to the menu; 0 exits. A passphrase mismatch prevents training.

Stop when something is not understood. An unanswered confirmation is not consent. A failed native command is retained in the session's journal/ and shown as a refusal, never translated into a successful lesson. The journal is readable local teaching history, not an independent proof receipt.

5.2 What your first files mean

For composition, open compose/RUN/report.html, then result.json and composition.successor.json. The first rehearsal should retain a predictive evaluation, readable model/evidence, a separately created trust journal and a local proposal. Its rehearsal: true state prevents preferred selection. A proposal's cell action numbers mean **0 idle**, **1 produce**, **2 service**, as hypothetical classifications only.

The companion's classifier and world lessons import **synthetic** evidence with policy new; that machine value means newly supplied input to this local lineage, not genuine customer observations or independently fresh data. The composition lesson uses policy rehearsal. Origin, evidence policy and qualification must be read separately.

Keep the original archive, product, practice records, passphrase and operator journal distinct. Only the private signing key is encrypted. A colleague can read model and evidence content without that passphrase.

5.3 Return to an existing session

The companion prints the exact command for your session. For the example layout:

```
| python3 START_PRACTICE.py --resume ../omega-practice
```

It reopens completed objects rather than training again. It refuses changed prepared inputs; use the native wizard in a new research branch to configure another mission. A missing or partial RUN/ is not a completed checkpoint: preserve it and start a new lesson directory after inspection.

The context generated for a proposal is explicitly synthetic and receives a provenance record. A new context does not extend the signed composition's expiry. If the composition itself is expired, create a new deliberate lesson rather than altering the clock or signed file. Real stale observations must be reacquired, never relabeled current.

5.4 The native route remains available

From the **product** folder, the original generator and wizard still work:

```
| node 07_NATIVE/examples/r9/make_example.mjs ../r9-demo
| node 07_NATIVE/start.mjs compose wizard
```

Choose language, form, parent NONE, mission ID guide-first-session, title Synthetic shared-resource lesson, organization Practice owner, objective Inspect local proposals only, and evidence context Explicit synthetic teaching data; not a factory. Choose train-neural, rehearsal, rights yes, and evaluation retirement no.

Use source ../r9-demo/formation.json, source label Synthetic tutorial, version 1, and resource contract ../r9-demo/RESOURCE_CONTRACT.json. Retain the displayed population and UTC expiry, allocator ratio, horizon 3, future delivery ZERO, exponent 1, beam 8. Use new paths for the workspace, confirmed plan and result. Inspect preflight before agreeing. The first-session cards provide these answers together.

The completion test: explain which evidence was used, what was actually fitted, which result remains unknown, what is prohibited, and how to reopen the exact object. A beautiful chart alone is not completion.

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. [3, 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 the guide example it is 1,000 utility points, not money.

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.

7 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. [6]

7.1 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:

```
evidenceStrength,outcomeGap,interventionFit,stakeholderRisk,
reversibility,driftRisk,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.

7.2 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.

7.3 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 is 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.

8 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. [3, 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 stay unknown. 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.

8.1 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 new companion does not collapse these states into one "success" badge.

9 Create a successor deliberately




Figure 2. The retained local candidate is not mutated while the learning branch is formed.

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.

Worked negative outcomes are included. The new guide tests an explicit classifier reuse branch that lacks eligible new evaluation and leaves the parent unchanged. Its record is not promoted to a new proof. The source product tests also exercise incompatible, inferior and interrupted paths.

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.

10 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. [6]

```
| 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.

10.1 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.

10.2 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.

11 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.

12 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
```

12.1 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.

12.2 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.

12.3 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.

13 Pin, inspect and propose

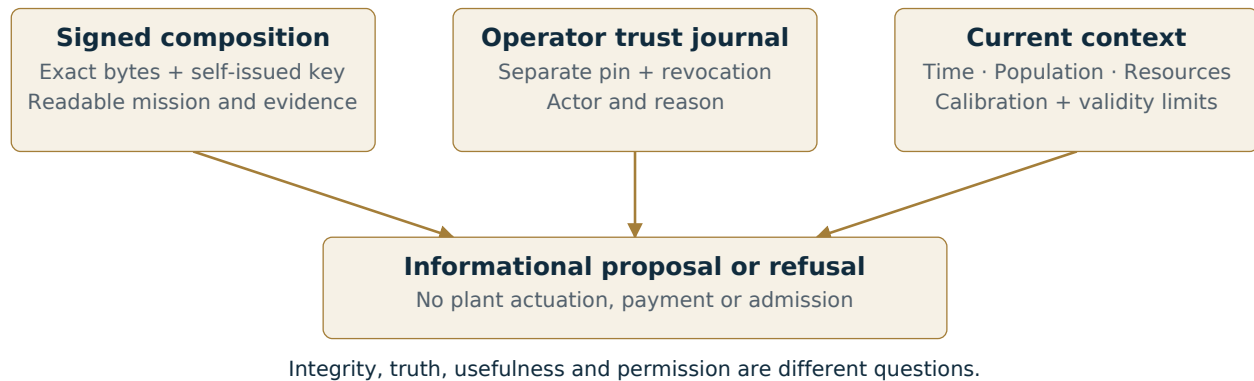


Figure 3. The public composition path checks separate records; a signature does not establish truth.

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.

14 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` 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.

15 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. [3, 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.

15.1 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.

15.2 Keep the strongest component, not the preferred story

ORCHESTRA's historical R8 study found that a symbolic successor predicted better yet lost the operating comparison. The customer retained a neural predictor and changed the allocator. Its later ordinary-regime gain did not survive the power-shortage stress. These are lessons about experimental discipline, not R9 savings or a promise that your mission will reproduce the same outcome. The complete historical manuscripts are retained in R9's publication bonus folder. [10]

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.

16 A practical operating cadence

This is a proposed cadence, not an automatic R9 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. R9 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.

17 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. [3, 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
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

R9 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.

18 Restore, move and protect your work

18.1 What the passphrase protects

Only the private signing key is encrypted. Mission descriptions, evidence, model parameters, reports and most metadata are readable. Keep those files in storage appropriate to their sensitivity. The bundled teaching objects use public QA fixture secrets and are read-only examples: never reuse their keys or trust pins for your work. A new personal practice session generates its own keys with the passphrase you enter. A colleague may be able to read a mission without knowing its passphrase. Losing the passphrase does not necessarily prevent public signature checking, but it can prevent key-possession and authorized-continuation operations. There is no recovery mechanism promised in the product. [6]

18.2 Use the matching export family

```
node 07_NATIVE/start.mjs export --file RUN/mission.successor.json --out CLASSIFIER_EXPORT
node 07_NATIVE/start.mjs restore --dir CLASSIFIER_EXPORT --out CLASSIFIER_RESTORE
node 07_NATIVE/start.mjs world export --file WORLD_RUN/world.successor.json --out WORLD_EXPORT
node 07_NATIVE/start.mjs world restore --dir WORLD_EXPORT --out WORLD_RESTORE
node 07_NATIVE/start.mjs compose export --file RUN/composition.successor.json --out
COMPOSITION_EXPORT
node 07_NATIVE/start.mjs compose restore --dir COMPOSITION_EXPORT --out COMPOSITION_RESTORE
```

Each destination must be new. Restore preserves the exact signed object and available history. It does not transfer the local trust pin, preferred pointer, external proof or production authority. Review and pin a composition again in the restored operator context.

Raw export restrictions, including available inherited restrictions for compositions, can cause refusal. Do not delete evidence manually while continuing to claim complete replayability. Full replay and a restricted disclosure are different deliverables.

18.3 Keep four things distinct

Preserve the original customer ZIP, the run/export, your secret and the operator trust journal. Back up according to your own approved retention policy. Test restoration before relying on it. Copying everything into one folder does not create organizational separation or protect against a compromised operator.

Keep the product directory unchanged and outputs outside it. Modifying product files breaks the known inventory and makes behavior/replay comparisons harder. A new product release should be extracted beside the old one, not written over historical experiments.

19 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 07_NATIVE/start.mjs	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 wizard</code> 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 / <code>compose help</code> ; <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.

What you see	Likely meaning	Correct next step
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.
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.

20 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. [6]

R9 language detail: menu 8 invokes the retained Academy entrypoint, 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.

21 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, $e_{k+1} \leq Le_k + \varepsilon$ gives $e_k \leq \varepsilon \sum_{j=0}^{k-1} L^j$. 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. [11]

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.

22 Advanced operator note: scarcity and strategic behavior

Conditional theory. R9's process cells are centrally configured assets, not independent strategic owners. The following mathematics identifies what would need a separate contract or implementation.

22.1 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

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

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. [12]

22.2 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. R9's ratio score is not a VCG mechanism; there are no such participant payments in the product. [13]

22.3 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. [3, pp. 17, 28–29]

23 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 R9 programs satisfy it. [14]

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 km² of collector area and 4,888 km² of total-system radiating surface. These are reference identities, not a design or a construction forecast. [15, 16]

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.

24 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.

25 Command reference and tested copies

Run native commands from the **customer root**. Run guide commands from the **guide root**. The complete one-line catalogue is in 01_READ/COMMANDS_EN.txt; the HTML edition has a copy button on each code block. Your viewer may wrap a printed line, so the catalogue is the preferred copy source.

25.1 Orientation and native menus

```
node 07_NATIVE/start.mjs help
node 07_NATIVE/start.mjs wizard
node 07_NATIVE/start.mjs world wizard
node 07_NATIVE/start.mjs compose wizard
node 07_NATIVE/legacy-start.mjs academy --language en
```

The last command is the documented Academy-only exception; do not use the legacy endpoint for formation or succession. For French, replace the final language value with `fr`.

25.2 Preflight, confirmation and formation

```
node 07_NATIVE/start.mjs preflight --plan PLAN.json
node 07_NATIVE/start.mjs confirm --plan PLAN.json --token FULL_TOKEN --out CONFIRMED.json
node 07_NATIVE/start.mjs create --plan CONFIRMED.json --out RUN
```

FULL_TOKEN is a placeholder: replace it with the complete token printed by preflight. Keep the confirmed plan beside its source plan when evidence paths are relative. Formation without a secret flag requests hidden terminal input. For the other families use world preflight, world confirm, world synthesize, or compose preflight, compose confirm, compose form. All destinations must be new.

25.3 Verification and return to a lesson

```
node 07_NATIVE/start.mjs verify --file RUN/mission.successor.json
node 07_NATIVE/start.mjs world verify --file WORLD_RUN/world.successor.json
node 07_NATIVE/start.mjs compose verify --file RUN/composition.successor.json
```

From the guide root, reopen a completed teaching session without repeating training:

```
python3 START_PRACTICE.py --resume ../omega-practice
```

25.4 Prepare all three examples without training

```
python3 08_TOOLS/prepare_workshops.py --product ../omega-product/SUCCESSOR_OMEGA --out
../omega-practice --language en
```

This lower-level preparer writes input files only. START_PRACTICE.py additionally writes its own session record; a folder made only by the preparer is not automatically a resumable companion session. For a noninteractive companion preparation, use START_PRACTICE.py --prepare-only together with explicit --product, --session and --language arguments.

25.5 Technical regression and automation

```
node --test --test-concurrency=1 07_NATIVE/tests/*.test.mjs
```

The wildcard command is a POSIX-shell example. For portable enumeration, the guide supplies 08_TOOLS/run_product_tests.py --product PRODUCT_ROOT --out NEW_REPORT_DIRECTORY; it passes explicit filenames without shell expansion. This is an optional broad test, not part of an ordinary first session.

Formation supports --passphrase-stdin or --secret-fd; authorized parent continuation uses --parent-secret-fd. Never place a passphrase itself in a command argument, environment variable, screenshot or support message. The companion obtains it through protected terminal input and never stores it. No tool in this guide grants an independent proof receipt or action authority.

26 Source, verification and distribution notes

This guide is by **Vincent Boucher, President, QUEBEC.AI & MONTREAL.AI**. It was prepared with AI assistance from the exact R9 archive, its source, its retained documentation, the supplied architectural white paper and the explicitly cited literature. The author line does not assert independent review or personal execution of every command.

The earlier guide's command checks were executed against a newly extracted, unchanged R9 customer ZIP. The original R9 acceptance report is preserved as historical release evidence; the guide's new execution report identifies only the paths and tests rerun here. Numerical synthetic exercises are not customer testimonials or new financial-benefit trials.

26.1 Retained qualification and new integration checks

The Guide 2.0 terminal journeys, companion sessions, Academy checks and negative-path records are preserved under `04_EVIDENCE_GUIDE_2_0/`; its earlier Guide 1.0 records remain inside that historical directory. Those sessions are not represented as newly rerun for this integration.

Current checks are recorded separately in `04_EVIDENCE/`: exact source identities; the unchanged R9 package inventory and 118 local regression tests; the new mission-design review and value-of-proof calculations; source-to-R9 references; and the revised manuscripts, commands, navigation and clean LaTeX rebuild. Read the acceptance record for the exact scope of each check. Documentary completeness is not authenticated evidence or permission to proceed.

Guide 2.0 corrected an inline formatter that suppressed spaces in commands such as `py -3` and `node 07_NATIVE/start.mjs wizard`. That correction, the complete command blocks and the literal text/HTML copies are retained. This edition adds the SEIZE and Gym preparation pathway. Navigation checks compare intended headings, not just valid page numbers. Source builds and rendered-page inspection remain separate from product tests.

What is not established: a controlled study with novice human participants, a new complete R9 release acceptance, current-patch runtime qualification, Windows/macOS execution, real Docker, full native-browser operation, external proof or measured customer benefit. These are not converted into passing labels by adding a guide.

The corresponding source paths are `07_NATIVE/start.mjs`, `lifecycle/wizard.mjs`, `lifecycle/evidence.mjs`, `world/service.mjs`, `world/planner.mjs`, `composition/cli.mjs`, `composition/contract.mjs`, `composition/service.mjs`, `composition/wizard.mjs`, `academy.mjs`, `examples/r9/make_example.mjs`, and the three R9 guides in `08_GUIDES/`. The source map records their hashes.

The PDFs, reading editions and LaTeX share the same authored text. Navigation checks distinguish a destination that merely exists from the page containing the intended heading. The acceptance report records actual checks and remaining host limits. No PDF/UA or assistive-technology certification is asserted.

The conceptual cover crop is generated artwork. Generated labels and promotional charts from the unused full image are excluded; the picture is not a facility or capability measurement. All guide diagrams are separately authored explanations, and all worked numerical examples state their assumptions. No font files, runtime binaries, private customer evidence or external credentials are distributed.

The complete bundle contains licensed commercial software. A separate publication-only archive excludes that software, the three architectural source PDFs and the practice runtime. It retains the new

guide-side mission worksheets and their document-review tool. Sharing the guide does not alter the original platform's licence, support, pricing, ownership assignments or customer rights. The guide is a companion to R9—not a replacement product release.

27 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.

28 Before formation: use SEIZE to choose the next decision

Source-derived operating discipline; implemented here as guide-side worksheets, not a new R9 engine. 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?** [4, pp. 20–23, 5, 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 R9 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.



Figure 4. A source-derived pathway with an explicit local boundary. SEIZE worksheets prepare a decision; the existing R9 workbenches execute bounded local research. Independent proof and admission remain separate.

28.1 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. [4, pp. 21–22, 5, p. 24]

First practical use. Open 09_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; R9 does not supply a live market radar that authenticates those options.

28.2 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. 09_MISSION_KIT/vop_example.json 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 18.4, leaving a net assumed value of 13.4. 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. [4, pp. 10–11, 15–16, 5, 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.

29 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. [5, 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.

29.1 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?
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. R9'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. [5, pp. 5–9]

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

Source family	Intended role	What an R9 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.

Source family	Intended role	What an R9 user can honestly retain now
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. GYM_MATURITY.md supplies the exact source stages with evidence questions. It deliberately assigns **no automatic G-level to R9 or a worksheet**. Showing roles does not establish independent custody. [5, pp. 8–9]

29.3 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. [5, 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 R9 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.

30 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. [5, 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.

No.	Source job	Required retained object
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. R9 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-R9 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. `JOB_CONCORDANCE.json` preserves that numbering and wording without merging the underlying accountable decisions. The checklist above is explicitly the Gym version. [4, pp. 29–30, 5, pp. 61–62] A work-and-proof job is not an automatically dispatched agent task, blockchain job or paid transaction.

31 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.

31.1 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. [4, pp. 14–15]

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

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

The reserve is not estimated merely by having a simulator. [5, pp. 21–22, 39–40]

The later white paper uses the lower-bound comparison and deductions retained by the R9 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. [3, 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 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.

31.2 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. [5, 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 R9 evaluator or a historical receipt to resolve the source difference silently.

31.3 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. [4, pp. 20–24, 5, pp. 29–32, 3, 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. R9 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. [4, pp. 1, 5–7, 5, pp. 2, 14–16]

31.4 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 R9 product. This integration introduces no wallet connection or balance check. [4, pp. 36–37]

The attached commercial ladders contain illustrative engagement prices and proposed services. They do not amend R9’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. [5, pp. 52–55, 62–64]

32 A worked mission: decide first, then use the supported route

Explicitly fictional teaching example, not a new benefit trial. A property operator wants to prioritize already identified supplier-invoice exceptions for human review. The first choice is not whether to let an agent authorize payments. It is whether a documented six-signal triage mapping warrants a bounded classification experiment.

The Gym paper’s Aurelia cycle includes a synthetic invoice-integrity environment, comparisons, demonstration admission, impairment and a further successor. Those are source-reported, synthetic and unaudited results. They are not R9 execution records, production authority or money recovered for this example. [5, pp. 44–47]

32.1 The owner changes the decision, not the label

Step	Filled teaching decision
Succession event	Reviewer backlog motivates testing a more consistent triage method; no measured loss is asserted.
Alternatives	Retain manual review; repair the rubric; compare an existing alternative; train a bounded candidate; reserve full automation.
Anti-recommendation	Poor source-to-signal mapping can make a well-fitted classifier useless. Audit the mapping before more training.

Step	Filled teaching decision
Experiment Zero	Review a small, deliberately selected sample of exception records for mapping consistency and decision-time availability. Sample size is an exercise input, not a statistical guarantee.
Gym representation	Identified exception, six documented governance signals, proposed label, later reviewer outcome, separate episode/group identities where relevant.
Hard stop	Unclear use rights, missing critical fields, incompatible adapter, hidden review work, or unauditable baseline.
R9 route	Classification only. Document extraction, fraud adjudication and invoice payment stay outside this adapter.
Next decision	Inspect the mapping audit. Proceed to a separately confirmed local plan, repair the mapping, retain the incumbent or stop.

The filled `MISSION_EXAMPLE.json` describes an **explicit synthetic rehearsal** with equally explicit comparator plans. It is a planning record, not R9 training data. Its review result can be `DOCUMENT_COMPLETE_FOR_MANUAL_REVIEW`; that code grants no permission and establishes no evidence truth.

From the guide directory:

```
python3 START_MISSION.py example --out ../mission-example.json
python3 START_MISSION.py review --file ../mission-example.json --out ../mission-review
```

The local review writes JSON, Markdown and escaped HTML. It checks declared fields, supported route choice, explicit evidence policy, missing comparator rationale, hard-stop descriptions, source references and separation of proof. It does **not** authenticate rights, inspect protected cases, estimate performance, select a winner or authorize formation.

A blank template is intentionally incomplete:

```
python3 START_MISSION.py template --out ../my-mission.json
```

After human review, start the existing practice companion for a synthetic exercise or the native R9 wizard for a real supported research plan. The tools are deliberately not wired to auto-confirm each other. The readable mission worksheet can assist entry of title, objective, baseline, horizon, constraints and evidence context; R9 still resolves and confirms its own authoritative execution plan.

32.2 Useful outcomes include not building

The included examples cover missing rights, an absent Beta, an unsupported document-to-agent request and a declared reserve option. An absent Beta leaves the limited incumbent comparison available where supported; it makes the stronger Beta-relative conclusion unavailable. Unknown rights or an unsupported mapping are reasons to stop before training. Preserve the decision and its reason rather than turning every completed worksheet into a candidate.

33 Preserve failure, substitute carefully, and re-underwrite

The Gym paper's **Negative Capability Graph** records rejected architectures, reward exploits, rights failures, verifier weaknesses, unsafe paths and non-transferable improvements, linked through shared causes, components and repairs. This is richer than deleting an unsuccessful run. [5, pp. 34–36]

The guide supplies a graph worksheet with nodes for the failure, affected component, supporting event, scope, corrective hypothesis and retest condition. It is guide-side planning memory; it is not a claim that R9 has gained a native cross-mission graph database. A refusal about stale evidence belongs in this graph as an actual refusal, not as evidence that the underlying model is bad.

Evaluate two changes separately. A revised candidate needs a new comparison. A revised Gym needs its own version, provenance and validation. Changing both and measuring only a better score can conceal a weakened test. Keep the old failed case for regression checking without calling it fresh final evidence.

The source's mission-gauge idea concerns substitution that preserves declared mission observables within preregistered tolerances while the evidence chain remains valid. It does not authorize a silent supplier switch or automatic proof inheritance. A real substitution drill should record exact components, compatible interfaces, behavior, costs, rights, uncertainty, restore state and new applicability. R9's local import/restore options and bounded interpreter agreement are useful evidence at their tested scope, not a proof of every external substitution. [5, pp. 16, 38–39]

33.1 A board cadence that can end in retention

The source proposes periodic and event-driven re-underwriting. The worksheet retains the event that would matter: new baseline capability, changed prices, expired rights, missing recovery, critical error, shifted population or increased proof burden. Every trigger names an owner, a due condition and a permitted local response. No background monitoring service is installed by this guide. [4, pp. 19–20, 40, 5, pp. 35–40]

Track whether a later cycle uses fewer hours, fewer repeated failures or lower proof burden **at comparable mission difficulty and proof quality**. The sources call mere reuse recurrence, not verified recursive improvement. Record the evidence for any stronger claim; no generation number proves compounding.

Sell the decision before the build. The attached productization sections support a readiness/SEIZE engagement that returns a go, hold, repair or stop decision before larger formation. Use that structure for a separately scoped proposal, not as an implied service included in the R9 purchase. A mission environment can be a strategically durable asset, but the source expressly distinguishes managerial valuation from financial-statement recognition. [4, pp. 46–47, 5, pp. 38, 52–55]

34 Integration register: what changed and where it came from

This is **User Guide 3.0, Mission Navigator & Gym Edition**, a companion to unchanged R9. It adds seven source-derived chapters, a bilingual mission-design kit, a local completeness reviewer and source-qualified accounting notes. The existing first-session companion, signed teaching records and original product remain unchanged. Historical Guide 2.0 acceptance is retained under 04_EVIDENCE_GUIDE_2_0/; current integration checks are under 04_EVIDENCE/.

Reading need	Authoritative source location	Guide use
Choose the reversible evidence path	Navigator §9, printed 20–22; PDF 30–32	SEIZE card and decision alternatives.
Separate institutional decisions	Navigator §10.1, printed 23; PDF 33	Five-decision role record.
Specify the environment	Gym §3, printed 5–9; PDF 17–21	Gym contract, five families and maturity questions.
Separate policy from institution	Gym §§2,4, printed 3–5,10–12	Route boundaries and component map.
Bound the 21 work orders	Gym Appendix C, printed 61–62; PDF 73–74	Exact job names and unassessed delivery checklist.
Preserve negative capability and real recursion	Gym §11, printed 34–36; PDF 46–48	Failure graph and renewal records.
Keep accounting source-qualified	Navigator §7.4; Gym §7.4; later WP §10	Protocol concordance and burden register.
Check access separately from authority	Navigator §16, printed 36–37; PDF 46–47	No new R9 wallet or authority rule.

The attached Navigator is **v5.0.0, 9 August 2026, 71 PDF pages**; the Gym is **v2.1.0, 10 August 2026, 83 PDF pages**. Their printed main-body counts are 61 and 71 respectively. The copies in 05_REFERENCE/ preserve the exact uploaded bytes; 04_EVIDENCE/INPUT_IDENTITIES.json records their original filenames and full hashes. Page links use physical PDF pages, while citations use printed pages.

The new machine-readable records use guide-specific schemas. They are not canonical product manifests, proof receipts or Authority Envelopes. The source crosswalk records what can be drafted now, what the bounded R9 route can do, and what requires separate people, environments and evidence. The original paper pseudocode is not executed as a security boundary.

What is not changed. No protected verifier, access gate, external provider, payment, live action, entitlement or production admission is added. The source papers are not imported as labeled training episodes. Their conceptual figures do not become test results. Neither their historical synthetic examples nor their dates supersede the exact product's command contract.

Best use of the complete package: rehearse safely with the existing first-session guide; use SEIZE to choose the real question; design a credible Gym and basis-risk record; run only a supported, explicitly confirmed local branch; retain both positive and negative evidence; commission separate proof only when it can change a consequential decision.

References

- [1] MONTREAL.AI / QUEBEC.AI. *SUCCESSOR Ω 19.0.0 R9: Customer Product*. Exact supplied customer ZIP: 43,940,303 bytes; SHA-256 0fa222c25ba7a52e2e28c05aaa706c11ed01cdbacdfdeaff263ab15fb5632872. Product unchanged by this guide. 2026.
- [2] MONTREAL.AI / QUEBEC.AI. *SUCCESSOR Ω 19.0.0 R9: Quick Start / Démarrage rapide*. Six-page supplied bilingual quick start; current source controls exact commands. 2026.
- [3] 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.
- [4] 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 R9 execution proof. Aug. 9, 2026.
- [5] 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.
- [6] MONTREAL.AI / QUEBEC.AI. *R9 native source, data contracts, composition guides and support terms*. Paths and hashes in 04_EVIDENCE_GUIDE_2_0/SOURCE_MAP.json. In particular 07_NATIVE/start.mjs, lifecycle, world, composition; 08_GUIDES and 04_LICENSE_AND_POLICIES_EN_FR. 2026.
- [7] Python Software Foundation. *Using Python on Windows*. Official launcher syntax; not a Windows qualification of this guide or R9. URL: <https://docs.python.org/3/using/windows.html> (visited on 09/06/2026).
- [8] Node.js Project. *Node.js Releases*. Runtime maintenance reference, not product qualification. URL: <https://nodejs.org/en/about/previous-releases> (visited on 09/06/2026).
- [9] Node.js Project. *Download Node.js*. Use an approved maintained LTS; record and test exact local version. URL: <https://nodejs.org/en/download> (visited on 09/06/2026).
- [10] V. Boucher. *SUCCESSOR Ω —ORCHESTRA Ω : Intelligence Under Shared Constraints*. Author edition 2.0, historical R8 study retained in R9's publication bonus. Not a new R9 economic evaluation. 2026.
- [11] 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/~jbraw/mpc/> (visited on 09/06/2026).
- [12] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004. URL: <https://web.s tanford.edu/~boyd/cvxbook/> (visited on 09/06/2026).
- [13] T. Roughgarden. *Incentives in Computer Science, Lecture 15: The VCG Mechanism*. Stanford University. Nov. 14, 2016. URL: <https://theory.stanford.edu/~tim/f16/l15.pdf> (visited on 09/06/2026).
- [14] 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 an R9 experiment. 2018. arXiv: 1804.04064.
- [15] 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).
- [16] 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).