

WORLD 8

From a persistent digital entity to a governed multi-society runtime — a tree-structured reconstruction of the research path, evidence, implementation, global position and next proof targets.

World 6 → World 7 → World 8

Company Society + Trading Society

Evidence-Bounded

Research & Architecture: Saeed Farokhi
Academic contact: Saeed.farokhi@ut.ac.ir
Evidence snapshot: 26 August 2026

Status: Z0-A design frozen; reference implementation started; Production NON-CLAIM

Purpose: reconstruct lineage, questions, evidence, global position, two-society model and next targets

World 8 in one tree

This dossier is not a marketing rewrite of the architecture. It is a reconstruction of what was asked, what changed under criticism, what exists in code, what remains open, and why the next milestone is a two-society executable proof rather than another larger document.

WORLD 8

- └ LINEAGE
 - | └ Entity Ark / World 6 – persistent entity, recovery-first operational core
 - | └ World 6.2 – resolution-native, portable brain, evidence-bounded candidate
 - | └ World 7 – genome model, proof registry, falsifiability, hard review
 - | └ World 8 / Z0-A – five-plane governed runtime; design frozen 24 Aug 2026
- └ RESEARCH QUESTION TREE
 - | └ identity: what remains the same when model/provider/session changes?
 - | └ truth: who owns canonical state if cognition is replaceable?
 - | └ authority: who may propose, approve, commit, execute and promote?
 - | └ effects: how do retries avoid duplicate real-world consequences?
 - | └ recovery: can the World be reconstructed on a clean host?
 - | └ evolution: can experience improve the system without self-authorized mutation?
 - | └ generality: can two radically different societies share one kernel?
- └ CURRENT EXECUTABLE EVIDENCE
 - | └ multi-provider farm prototype
 - | └ versioned Proforma Skill
 - | └ proposal-only Mason v0
 - | └ independent Evaluator v0
 - | └ MCP / canonical gateway + governance evidence
- └ SOCIETIES
 - | └ Company – Secretary first, Sales next
 - | └ Trading – Forecast → Portfolio → Risk → Intent → Effect → Reconciliation
- └ GLOBAL POSITION
 - | └ durable execution: mature prior art exists
 - | └ identity/authorization: active frontier and standards work exists
 - | └ self-improvement: active research exists
 - | └ World 8 hypothesis: integrate these under one governed, evidence-bound contract
- └ NEXT PROOF
 - | └ ONE KERNEL · TWO SOCIETIES · ONE CONFORMANCE SUITE · REPRODUCIBLE EVIDENCE

Claim discipline. Architecture, implementation and evidence are separate. The frozen Z0-A document explicitly states that local tests, diagrams, metaphors or elapsed time do not upgrade a claim. This report preserves the same rule. [P-W8]

Z0-A ARCHITECTURE

FROZEN

Final design baseline; not production.

REFERENCE RUNTIME

STARTED

Several real components are executable.

TWO-SOCIETY PROOF

OPEN

The decisive conformance experiment has not yet passed.

Documented lineage: 12 → 26 August 2026

The compressed timeline matters, but only as evidence of iteration velocity. It is not a substitute for production evidence. The documented archival lineage begins with the Entity Ark / World 6 operational standard dated 12 August and reaches a frozen World 8 baseline, early runtime components, and a second Society architecture by 26 August.

Date	Branch	What changed	Evidence boundary
12 Aug	Entity Ark / World 6	Operational core centered on stable Entity identity, PostgreSQL truth, policy-safe Brain Gateway, append-oriented audit, outbox/effect semantics, recovery, Human Root authority and E0-E5 evidence.	21 JS + 6 Python contract/local tests documented; clean-directory restore; fresh-machine PostgreSQL restore and live two-provider proof remained open. [P-W6]
20–21 Aug	World 6.2	Resolution-native candidate, portable brain packs, safe partial resolution, backward-compatible migration and larger candidate validation package.	Public DOI lineage exists. Historical candidate evidence is preserved but does not automatically upgrade later generations. [P-W62]
22–23 Aug	World 7 v0.2	Architecture became a proof-bearing research object: Universal Genome, authority algebra, Spine, Proof Registry, explicit gaps, evidence states, and a first phenotype.	18/18 seed reference tests, 4096/4096 bounded assertions, 2048 obligations with 1720 explicitly OPEN. Not production, not restore-proven. [P-W7]
24 Aug	World 8 / Z0-A	Metaphor was subordinated to five fixed planes, strict duties separation, Objective/Revision contracts, CAS/fencing, effect receipts, clean-host recovery gates and explicit non-claims.	FINAL DESIGN BASELINE / NOT PRODUCTION. [P-W8]
24–25 Aug	Reference implementation	Versioned Proforma Skill, proposal-only Mason, independent Evaluator, MCP gateway and Human Root/promotion evidence were committed.	Real code exists, but it is early and incomplete. [GH-SKILL][GH-MASON][GH-EVAL][GH-GATE]
26 Aug	Two-Society maturity baseline	Company and Trading were placed under one shared World; hierarchical learning and cross-society promotion were defined as implementation blueprint.	Company partially executable; Trading runtime still open; shared conformance proof not yet run.

Velocity, correctly stated: the project traversed several architecture/review/implementation cycles in roughly two weeks. That is unusual iteration speed under a very small operating budget, but it does not make its deployment maturity comparable to large production platforms.

The questions that drove the architecture

The project did not mature by adding features to a chatbot. It matured by repeatedly replacing vague ideas with narrower systems questions that could be falsified.

QUESTION TREE

└─ IDENTITY

- | └─ If the model changes, is it still the same entity?
- | └─ If the session disappears, what survives?
- | └─ Which IDs are canonical and which are only provider provenance?

└─ TRUTH & STATE

- | └─ Can cognition rewrite the truth it reasons about?
- | └─ What is canonical versus derived/projection state?
- | └─ How are stale writers, replay and tampering handled?

└─ AUTHORITY

- | └─ Does a Role name create permission? (Answer: no.)
- | └─ Who may propose, evaluate, promote and execute?
- | └─ How is Human Root bounded without becoming a bypass?

└─ EFFECTS

- | └─ What does 'send it' actually authorize?
- | └─ What happens when a provider times out after possibly executing?
- | └─ How do retries remain intent-bound and collision-safe?

└─ RECOVERY

- | └─ Is restart on the same host a recovery proof? (Answer: no.)
- | └─ Can identity/history/authority be reconstructed and followed by a fresh governed write?

└─ DEVELOPMENT

- | └─ Can a system learn from experience without self-authorizing changes?
- | └─ Can Observation be independent from the developer that consumes it?
- | └─ Can evaluator identity be proved rather than assumed?

└─ GENERALITY

- └─ Is the architecture just a Secretary architecture?
- └─ Can Company and Trading obey exactly the same World laws without a core fork?

How answers hardened over generations

World 6 answer

Keep identity/state/policy outside the Brain; use a World Registry, Event Ledger, Brain Gateway, policy gate, recovery plan and Human Root.

Important early rule: no new operational-core complexity without a real failure/need, Build-vs-Integrate analysis, repeatable test/falsifier, migration/rollback and Human Root promotion approval. [P-W6]

World 7 answer

Turn anatomy into explicit contracts and proof obligations. Distinguish capacity from authority, role from holder, canonical history from model proposals, and test counts from actual evidence level.

The Proof Registry made "OPEN" a first-class status rather than hiding missing evidence behind large test counts. [P-W7]

World 8 answer

Freeze a normative five-plane design. No LLM, UI, Mason, workflow engine, connector or repository branch may be canonical truth. Every accepted change crosses authority, policy, CAS, idempotency, append history, fencing and evidence gates.

Current answer

Do not broaden theory further. Build two very different Societies on one shared kernel, make both pass the same conformance laws, restore them on a clean host, and let evidence decide whether the abstraction generalizes.

The architecture at its current mature design baseline

The central Z0-A thesis is that cognition and interfaces are replaceable, while accountable identity, canonical state, authority, objectives and accepted history are not. [P-W8]

WORLD 8 / FIVE FIXED PLANES

- └─ CANONICAL SPINE
 - | └─ identity + canonical event history + head/version
 - | └─ active Objective Contract + Revision activation
 - | └─ authorization + promotion/rollback events
 - | └─ controlled append; no runtime rewrite of committed history
- └─ OPERATIONAL
 - | └─ Entity / Role / RoleBinding / Task
 - | └─ ingress + auth + routing + Brain Gateway + context
 - | └─ connectors + outbox dispatch coordination + read models
- └─ OBSERVATION
 - | └─ telemetry → normalize → window → predicate/SLO
 - | └─ attributed gap_signal (measurement, not canonical truth)
- └─ DEVELOPMENT / MASON
 - | └─ candidate revision only; no direct authority, evaluation or promotion
- └─ EVIDENCE / GOVERNANCE
 - | └─ Proof Registry + falsifiers + evidence refs
 - | └─ evaluator receipts + approval + policy
 - | └─ promotion gates + audit

Separation of duties

Normative chain: Observation measures → Mason proposes → Evaluator judges → Promotion Authority promotes/rejects → Canonical Spine records → Effect Executor executes only authorized obligations. Combining those powers without an independent boundary violates Z0-A. [P-W8]

Identity, Role and authority

An **Entity** has stable identity and lifecycle. A **Role** is a responsibility slot, not an identity. A **RoleBinding** binds an entity/holder/executor to a role under explicit scope, grants, version and effective interval. A provider change must not change canonical World/Entity/Society/Role/Task lineage.

Objective and Revision

Objectives cannot live only in prompts. Objective Contracts are immutable/versioned and bind objective, success metrics, hard constraints, authority ceiling, risk class, observation contract, allowed adaptation surface and promotion policy. Runtime changes are represented as parent-linked Phenotype Revisions with migration, rollback, risk and evidence references.

Canonical commit mechanics

The Spine requires commit-time compare-and-swap (CAS), fencing, append-oriented history and intent-bound idempotency. Two writers starting from the same head cannot both silently win. Reusing the same idempotency key with different semantic intent is a collision and must fail.

External effects

External actions are obligations, not hidden handler side effects. The design path is authorization → canonical commit/outbox → dispatch attempt → provider observation → reconciliation → Effect Receipt → settlement. Universal exactly-once is explicitly not claimed; the target is scoped effectively-once semantics with stronger treatment for non-idempotent destinations.

Recovery

Same-host restart is insufficient. Clean-host recovery must reconstruct schema/migrations, ledger, objectives, revisions, artifacts, bindings and checkpoints, verify lineage/head, obtain fresh fencing authority and complete a fresh governed write. Until that drill is evidenced, recovery remains OPEN.

How the project learned not to confuse activity with proof

World 7 deliberately made evidence status visible. This is central to understanding the project because its maturity came partly from finding weaknesses in its own proof story.

World 7 evidence item	Recorded result	What it does NOT prove
Seed reference tests	18/18 PASS	Not production deployment, security or clean-host continuity.
Bounded local assertions	4096/4096 PASS	Not 4096 independent end-to-end proofs.
Proof obligations	2048 total: 1720 OPEN; 132 DEFINED_E1; 160 PASS_E2_LOCAL_MODEL; 30 historical candidate; 6 PASS_E2_LOCAL_REFERENCE.	2048 obligations do not mean 2048 passed proofs.
World 6.2 ancestor evidence	Historical candidate evidence preserved	Does not automatically substitute for live World 7/8 integration proof.

Review-driven correction. Later review found some assertions too self-referential/vacuous and some mutations surviving. The response was not to defend the inflated story; it was to narrow the core, move “Living Genome” to exploratory lineage, strengthen replay/ idempotency/evaluator evidence, and stop claiming autonomous evolution without the required population/selection evidence.

The evidence ladder now implied by the project

EVIDENCE

- └─ E0 – idea / prose / unexecuted claim
- └─ E1 – defined contract / schema / testable predicate
- └─ E2 – local executable evidence
- └─ E3 – integration / external system / restore / adversarial evidence
- └─ E4 – limited real operational evidence
- └─ E5 – sustained production evidence

The project's present public claim ceiling is intentionally below production. The strongest next upgrades come from cross-system conformance, clean-host recovery, external effect reconciliation and limited operational evidence—not from generating more architecture pages.

What is actually implemented today

The frozen Z0-A PDF is a design baseline; the repository changed after that freeze.

The correct current description is therefore **Reference Implementation:**

STARTED / PARTIALLY EXECUTABLE.

EXECUTABLE / REPOSITORY EVIDENCE

- |— Multi-provider Farm prototype
 - | — recorded mixed provider outcomes without treating provider failure as identity failure
- |— secretary.proforma Skill
 - | — versioned executor
 - | — validation + arithmetic + PDF artifact generation
 - | — storage/execution evidence path
- |— Mason v0
 - | — reads independently produced gaps
 - | — proposal-only configuration patch
 - | — forbidden from identity/authority/objective/spine/promotion changes
- |— Evaluator v0
 - | — reads Mason proposal
 - | — constructs candidate and deterministic regression result
 - | — records PASS/FAIL without promotion
- |— MCP / World Gateway
 - | — world8_status + get_context
 - | — begin_turn + commit_turn
 - | — generate_proforma + record_skill_feedback
 - | — mason_propose_revision + evaluate_mason_proposal
- |— Governance evidence
 - | — Human Root / promotion-gate hardening and recorded acceptance artifacts

Multi-provider Farm

A recorded farm run produced successful responses from Groq, Mistral and Hugging Face/DeepSeek while the chosen OpenRouter and Gemini model routes failed. The significance is architectural rather than benchmark performance: a provider can be available, unavailable or deprecated without becoming the identity of the Entity. [GH-FARM]

Versioned Proforma Skill

The committed `secretary.proforma` executor includes deterministic normalization, positive-integer validation, line-item arithmetic, PDF generation, Persian rendering, object storage and execution lookup. It is concrete evidence that “Skill” can be more than a prompt. [GH-SKILL]

Mason and Evaluator

Mason's runtime code is explicitly proposal-only and allowlists the patch surface. The independent evaluator then reconstructs the candidate and runs deterministic regressions. These v0 components do not prove autonomous evolution, but they implement the architectural separation between proposal and judgment. [GH-MASON][GH-EVAL]

Gateway

The Grok MCP gateway exposes canonical context/turn tools and separately exposes Skill feedback, Mason proposal and evaluation tools. Its own instructions state that World 8 identity lives in the canonical Spine rather than Grok. [GH-GATE]

Not yet true: a fully production-grade provider-neutral runtime, a completed cross-society Mason, clean-host restore proof, live real-money Trading, or a generality claim proven by two production Societies.

Company Society: the first operational phenotype

Company is the first domain where the World is forced to interact with humans, documents, channels and business effects. Secretary is currently the most mature operational phenotype; Sales is the next obvious role/phenotype, but is less implemented.

WORLD 8

- └─ COMPANY SOCIETY
 - └─ Entity / Role topology
 - | └─ Secretary – Role; may be backed by its own Entity when lifecycle warrants
 - | └─ Sales Manager – Role / future independent Entity as justified
 - └─ Shared operational services
 - | └─ Task Bus + context + Brain Gateway
 - | └─ Telegram / email / forms / document channels
 - | └─ Memory + artifacts + evidence
 - | └─ effect authorization + outbox + receipt
 - └─ Current executable Skill
 - | └─ secretary.proforma
 - └─ Development loop
 - | └─ feedback → observation/gap → Mason proposal → Evaluator → governed promotion

Why Secretary is a useful first phenotype

It is not a toy simply because it generates documents. A Secretary forces the architecture to bind identity, customer/task context, deterministic financial calculations, artifact generation, external delivery, corrections, audit, provider provenance and long-lived organizational memory.

Entity inflation is intentionally avoided

A Role is not automatically promoted into an Entity. Secretary or Sales becomes a distinct Entity only when independent lifecycle, memory, budget, queue, freeze/repair boundary or accountability justifies it. This keeps the taxonomy stable and prevents “every function is an agent” architecture.

Company proof target

A minimal Company conformance path is: Task → RoleBinding → Brain Proposal → governed decision/authorization → deterministic Skill/effect → Receipt → Evidence. A provider switch must leave canonical identity and task lineage unchanged.

Trading Society: the deliberate stress test

Trading was chosen because it is not another enterprise workflow demo. It introduces probabilistic claims, fast-changing data, capital, risk ceilings, concurrency, ambiguous external effects and reconciliation. If the same World laws survive this domain, the generality hypothesis becomes meaningfully harder to dismiss.

TRADING SOCIETY

- └─ Market Data Fabric (outside Canonical Spine)
 - └─ raw/high-volume data retained selectively; exact decision inputs become evidence artifacts
- └─ Forecast Hall
 - └─ independent Forecaster roles / Strategy Skills
 - └─ immutable raw Forecasts
 - └─ outcome_contract_ref for falsifiability
 - └─ calibration + proper scoring
 - └─ dependency-aware aggregation; no majority vote authority
- └─ Portfolio
 - └─ proposal under Objective Contract and current portfolio/capital head
- └─ Risk
 - └─ NARROW_ONLY relative to Objective Contract
 - └─ veto / approved ceiling
 - └─ never silently widens risk or mutates the proposal
- └─ Trade Intent Compiler
 - └─ pure deterministic transformation; authority-free; no reservation inside compiler
- └─ Governed acceptance
 - └─ accept_intent + reserve_capital + reserve_risk_budget + advance_head atomically
- └─ Effect path
 - └─ authorization → outbox → broker → receipt → reconciliation
- └─ Evaluation
 - └─ Forecast quality
 - └─ Decision quality
 - └─ Risk/Allocation attribution
 - └─ Execution quality

Critical Trading decisions already closed

Area	Decision	Why it matters
Market data	World owns decision knowledge/evidence, not all raw history. Market Data Fabric is separate from the Spine.	Prevents canonical-state explosion and lets retention be earned by evidence.
Forecast	Forecast $\neq$ Decision $\neq$ Order. Raw forecasts are immutable; recalibration creates a separate object.	Prevents retrospective rewriting and authority leakage from prediction to execution.
Outcome	Every claim needs a versioned outcome contract: variable, return convention, source, adjustment, evaluation window.	Makes forecasts falsifiable.
Risk	Risk Policy is NARROW_ONLY relative to Objective Contract.	Operational risk processes can tighten but not mint broader constitutional authority.
Intent	Trade Intent Compiler is pure/deterministic and authority-free.	Separates decision materialization from capital/risk state change.
Reservation	Capital and risk capacity are reserved atomically at governed intent acceptance, not inside compiler.	Prevents double-spending capacity and race conditions.
Ambiguity	Possible-delivery timeout triggers reconciliation before retry; ambiguous reservation remains held.	Prevents duplicate orders and premature risk release.
Emergency	Human Root cannot bypass intent/outbox/audit. Emergency capability is narrow/reduce-only.	Keeps constitutional guarantees during stress.

What is NOT implemented yet

The Trading Society is architecturally mature but has no equivalent end-to-end runtime to Secretary yet. The immediate target is paper execution, not real money: deterministic market input → forecast → risk decision → exact intent → paper broker → receipt → reconciliation.

Entity → Society → World learning without domain leakage

The latest implementation blueprint adds a three-scope learning model. It is an extension built after the frozen Z0-A baseline; it should be recorded as a new ADR/contract rather than retroactively attributed to the 24 August PDF.

EXPERIENCE

- | ENTITY scope
 - | └ e.g. Secretary formatting / one Forecaster calibration behavior
- | SOCIETY scope
 - | └ e.g. Company lead workflow / Trading market-microstructure logic
- | WORLD scope
 - | └ e.g. context loss after provider failover
 - | └ e.g. idempotency / credential isolation / migration pattern
 - | └ candidate enters World only after evidence of domain independence

COMMON DEVELOPMENT PATH

```
Experience → Observation → attributed gap → Mason candidate → independent Evaluator
      → Conformance across affected scopes → Promotion Authority → canonical revision
```

Mason is a Role/Capability before it is an Entity

Company Mason, Trading Mason and World Mason can be separate RoleBindings using common framework code while maintaining distinct scopes, memories and authority. Mason becomes an independent Entity only if its lifecycle, memory, budget, crash isolation and accountability actually require one.

A correction to the naive promotion rule

Cross-society recurrence is strong evidence of domain independence, but “exactly two Societies” should not be a constitutional constant. The real gate is evidence: scope attribution, affected contracts, conformance tests, evaluator independence and promotion authority. Two independent Societies are a powerful initial test, not a magical number.

Implementation principle: Shared evolution engine; isolated society experience; evidence-gated cross-society promotion.

How generality will be tested rather than asserted

The strongest repository strategy during the proof phase is one monorepo with a shared Core and two reference Societies. A separate Trading architecture would make divergence easy and generality hard to falsify.

```
world8/
├─ architecture/
├─ contracts/
├─ runtime/
├─ conformance/
│   ├─ cases/
│   ├─ adapters/company/
│   └─ adapters/trading/
├─ societies/
│   ├─ company/
│   └─ trading/
├─ apps/world-control-center/
├─ tests/
└─ evidence/
```

Core-entry rule

A capability belongs in Core only when it is domain-independent and its owner plane, writer authority, idempotency, rollback, falsifier, evidence level and restore behavior are known. Market-only logic remains in Trading; company-only logic remains in Company.

Proposed common conformance laws

Law	Company scenario	Trading scenario
Provider-independent identity	Switch Secretary brain while task/history remain same.	Switch forecast brain while forecast/entity lineage remains same.
Proposal ≠ Authority	LLM cannot send document by itself.	Forecast/ensemble cannot place order by itself.
CAS stale-write rejection	Two task updates from same head: one stale-rejects.	Two intents from same portfolio head: stale path recomputes.

Law	Company scenario	Trading scenario
Intent-bound idempotency	Same document effect retry deduplicates; changed payload requires new authorization.	Same order intent retry deduplicates; changed qty/price creates new intent.
Observation noncanonical	Telemetry cannot mutate Company truth.	Market telemetry/metrics cannot mint portfolio truth.
Mason no self-promotion	Secretary Skill proposal cannot activate itself.	Trading strategy/runtime proposal cannot activate itself.
Clean-host restore	Restore Company lineage then make fresh governed write.	Restore then reconcile broker before fresh trading capability.

Future generality claim: not “two dashboards exist,” but “two radically different domains import the same kernel, pass the same invariant suite, recover under the same truth/authority rules, and produce inspectable evidence.”

Where World 8 sits relative to the 2026 agent infrastructure frontier

The credible positioning is not “nobody else has thought of this.” In 2026, multiple major vendors and standards efforts are converging on durability, persistent agent state, identity, authorization, audit and agent evaluation. World 8 is interesting because it attempts to make these concerns one governed architecture rather than because it invented every primitive.

External work	What it already does well	Implication for World 8
OpenAI Agents SDK [EXT-OAI]	Agent harness, sandbox separation, state externalization, snapshotting/rehydration and fresh-container continuation.	Durability and external state are not novel by themselves. World 8 must differentiate at identity/authority/evidence semantics and multi-society governance.
LangGraph [EXT-LG]	Built-in checkpoint persistence, threads, memory, human-in-the-loop, time travel and fault-tolerant resume.	World 8 should integrate mature runtime primitives where possible rather than reimplement checkpoint engines.
Microsoft Agent Framework [EXT-MS]	Durable sessions, checkpointed workflows, distributed hosting, persistent threads, reliable orchestration.	Microsoft is far ahead in production hosting maturity. World 8’s current advantage cannot be scale or deployment breadth.
Temporal [EXT-TEMP]	General-purpose crash-proof durable execution, replay, long-running workflows, retries, observability and AI-agent use cases.	Durable workflow mechanics should be considered an integration substrate, not a novelty claim.
Anthropic agent engineering [EXT-ANT]	Rigorous agent evals, evaluator/optimizer patterns, simplicity and complexity discipline.	Independent evaluation is established practice; World 8 must prove that its evaluator identity/promotion/evidence chain adds meaningful governance value.
NIST agent identity work [EXT-NIST]	Formalizes software/AI agent identity, authorization, auditing and non-repudiation as a current cybersecurity problem.	Confirms World 8 is addressing a real frontier problem; it also invalidates “nobody else is doing identity” claims.
IETF NHE architecture [EXT-IETF]	Persistent identity across sessions/hosts, bounded authority, tamper-evident records, operator control.	There is strong conceptual overlap. World 8 must compare schemas/invariants precisely and avoid uniqueness claims without exhaustive prior-art review.
AgentDevel / self-improvement		Mason/Evaluator alone is not novelty; the question is whether its placement inside a

External work	What it already does well	Implication for World 8
research [EXT-AD][EXT-SI]	Externalized improvement, release engineering, evaluation gates, controlled evolution from experience.	broader authority/effect/recovery architecture is useful.

Defensible differentiation hypothesis

World 8’s potentially distinctive object is the **integration contract**: provider-independent Entity identity + canonical truth outside cognition + versioned objectives + RoleBinding authority + CAS/fencing + intent-bound effects/outbox/reconciliation + independent Observation + proposal-only Development + independently attributable Evaluation + governed Promotion + claim-specific Evidence + clean-host restore + multiple Societies under the same laws.

Scientific wording: this is an *architectural differentiation hypothesis*, not established novelty or superiority. The next evidence must show that the integrated abstraction removes real failure modes or reduces engineering complexity in multiple domains.

Why the development velocity is notable — and how not to overstate it

Development to date has been founder-led with a very small operating budget and without a conventional staffed engineering or research organization. AI systems, cloud services and existing open-source infrastructure have been used as force multipliers. That context makes the iteration velocity notable, but it does not make the project operationally equivalent to platforms backed by large dedicated teams.

CURRENT RESOURCE PROFILE

Lean

No large institutional lab or conventional engineering organization behind the build; limited operational spend.

DEVELOPMENT MULTIPLIER

AI-heavy

Multiple models/providers used for research, critique, coding and evaluation support.

STRENGTH

Iteration

Fast architecture → review → implementation feedback cycles.

RISK

Bus factor

Architecture knowledge, implementation, operations and commercialization are too concentrated.

What additional team capacity changes

The project is now at the point where extra skilled collaborators can increase evidence quality more than additional conceptual expansion. The highest-value contribution is not simply more coding capacity; it is independent ownership of specialized proof surfaces.

COLLABORATION TREE

- └─ Distributed systems / database
 - | └─ CAS/fencing, event sourcing, transactional outbox, migrations, restore drills
- └─ Security / identity / authorization
 - | └─ principal binding, credential isolation, threat model, cryptographic attestations
- └─ AI runtime / providers
 - | └─ Brain Gateway adapters, provider-neutral contracts, context compilation, evals
- └─ Product / UI / UX
 - | └─ World Control Center, society plugins, evidence/authority visibility
- └─ Quant / trading engineering
 - | └─ data fabric, forecasting, calibration, risk, paper broker, reconciliation
- └─ DevOps / SRE
 - | └─ clean-host deployment, observability, chaos/load/recovery evidence
- └─ Applied research
 - | └─ conformance design, falsifiers, mutation testing, prior-art comparison
- └─ Academic / standards review
 - | └─ NIST/IETF alignment, reproducibility, publication-quality critique
- └─ Commercial / pilots
 - | └─ design partners, narrow paid wedge, regulated/enterprise use cases

Collaboration stance: the project is open to technical, scientific, product, infrastructure, data, security, quant, standards and design-partner collaboration. The immediate objective is to distribute responsibility around falsifiable milestones, not to inflate headcount.

Where this architecture is useful — and where it is not

World 8 is intentionally over-engineered for a simple one-shot chatbot. Its value rises when identity, authority, history, recovery, external consequences and long-lived organizational context matter.

Use case family	Why World 8 may help	Relevant proof requirement
Persistent digital workers	Brain/provider replacement without losing accountable identity or task history.	Provider handoff + role/task continuity.
Enterprise document operations	Versioned Skills, explicit authority, artifacts, receipts, correction/evidence loop.	Low-risk real effect + reconciliation + audit.
Sales/procurement/supplier automation	Long-lived entity memory plus bounded external communication/effects.	Recipient/payload-bound authorization and idempotency.
Financial/trading agents	Strict separation of probabilistic prediction, risk authority, exact intent and broker effect.	Paper/live evidence, OOS evaluation, reconciliation, risk containment.
Regulated/high-risk workflows	Explicit objective/policy versions, least authority, evidence receipts, audit and restore.	Independent security/compliance review; stronger evidence level.
Multi-model organizations	Model providers become replaceable cognitive resources rather than owners of the entity.	Multi-provider operational continuity.
Agent factories	Shared World laws with Society-specific phenotype/contracts.	Repeated conformance across independently different Societies.

Not a target: “AGI”, consciousness, biological life, guaranteed profitable trading, universal exactly-once effects, complete security or production readiness by declaration. Z0-A explicitly excludes these claims. [P-W8]

The UI/UX as an architectural proof surface

The shared Control Center is not meant to hide complexity behind a dashboard. Its job is to make architecture inspectable: one World, common services, two Societies, and explicit paths from experience to governed change.

WORLD 8 CONTROL CENTER

- |— WORLD SERVICES
 - | |— Identity / Registry / Spine
 - | |— Objectives / Roles / Bindings / Tasks / Skills / Brains
 - | |— Effects / Receipts / Evidence / Observation
 - | |— Mason / Evaluator / Governance / Recovery
- |— COMPANY VIEW
 - | |— Secretary
 - | |— Sales
 - | |— Company-specific workflow panels
- |— TRADING VIEW
 - | |— Forecast Hall
 - | |— Portfolio / Capital / Risk
 - | |— Execution / Reconciliation / Attribution

A candidate revision should be inspectable end-to-end: originating incident/experience → attributed gap → candidate scope → proposed diff → evaluator result → conformance outcomes → promotion/rejection → effect on subsequent metrics. This turns the UI into a governance and evidence interface, not merely an operational console.

The strongest objections the project must survive

Risk / objection	Why it is real	Required response
Architecture is larger than implementation	Most of the design surface is not yet end-to-end executable.	Freeze theory; implement minimum reference paths and conformance before new abstractions.
Reinventing solved infrastructure	Temporal, LangGraph, Microsoft and database ecosystems already solve major durability pieces.	Integrate mature components where they satisfy invariants; reserve custom code for World-specific semantics.
Novelty overclaim	Identity, authorization, durability, evaluation and self-improvement all have active prior art.	Claim integrated hypothesis, not primitive novelty; produce a formal related-work matrix.
Evaluator gaming / self-referential evidence	An adaptive system can optimize to its evaluator; World 7 already exposed weak proof families.	Independent datasets/falsifiers, mutation gates, external evaluation and explicit evidence lineage.
Security boundary leakage	Service keys, connectors and model-generated code can collapse intended duties separation.	Credential separation, least privilege, adversarial tests and independent security review.
Trading backtest illusion	Backtests can overfit; economic performance is path-dependent and cost-sensitive.	Predefined benchmarks, OOS/walk-forward, transaction costs, calibration, limited-live canary only after evidence gates.
Documentation drift	Post-Z0-A Trading/learning decisions can diverge from the frozen PDF.	Convert them into ADRs/contracts; never silently rewrite archival baseline.
Founder concentration / low bus factor	Knowledge and execution remain concentrated.	Independent maintainers, conformance docs, reproducible setup, evidence ownership and code review.

Z0-A architectural failure criteria already anticipate this discipline: provider/session becoming canonical identity, effects bypassing authorization/outbox, stale writers overwriting, proposer/evaluator collapse, restore being only same-host restart, or OPEN becoming PASS merely from test volume are architecture failures—not acceptable shortcuts. [P-W8]

Forward engineering targets

The dates below are planning targets, not guaranteed outcomes. The next report should state PASS / OPEN / FALSIFIED for each target instead of retroactively redefining success.

Target	Engineering milestone	Evidence required
31 Aug 2026	Clean world8 repository + contracts + top-level conformance skeleton.	Old lineage preserved; Company and Trading stubs consume one shared kernel; no secrets/raw licensed data in Git.
5 Sep	Company minimum reference path.	Secretary end-to-end effect + receipt + evidence; minimal Sales surface can remain a stub.
12 Sep	Trading paper reference path.	Deterministic input → Forecast → RiskDecision → TradeIntent → paper broker → receipt → reconciliation.
18 Sep	Two-Society Conformance run.	Same invariant suite executes against Company and Trading adapters with explicit PASS/FAIL report.
22 Sep	Clean-host restore drill.	Fresh environment reconstruction + verification + fresh fenced governed write; broker reconciliation before Trading resumes.
30 Sep	World 8 Executable Two-Society Reference v0.1 + Control Center.	Reproducible package, shared UI, evidence registry and honest claim status.

The ten proofs that matter more than the next 10,000 lines of prose

NEXT PROOF SET

- └ provider switch preserves canonical identity
- └ stale writer cannot overwrite
- └ same intent retry cannot duplicate economic/business effect
- └ changed payload cannot reuse prior authorization
- └ Mason cannot self-promote
- └ evaluator identity/result is independently attributable
- └ Company and Trading import the same kernel
- └ both pass the same conformance laws
- └ clean-host restore ends with a fresh governed write
- └ at least one low-risk external effect completes receipt + reconciliation

End-state of the next report: it should not ask a reader to imagine the architecture. It should let the reader inspect what ran, what failed, what recovered, and which claims legitimately moved upward.

Technical assessment at 26 August 2026

World 8 is neither a finished product nor merely a speculative paper. It is best described as a mature design baseline with an early executable reference path and a deliberately hard next falsification target.

DESIGN MATURITY

High

Z0-A is coherent enough to freeze as the current architecture and stop unconstrained conceptual expansion.

IMPLEMENTATION MATURITY

Early

Several important components run; the common kernel and full recovery/conformance story are incomplete.

SCIENTIFIC MATURITY

Improving

The project has learned to reduce claims after review and separate evidence status from test volume.

COMMERCIAL MATURITY

Pre-wedge

A Company reference path exists, but the first external paid design partner and narrow wedge remain to be proven.

Current engineering estimates — not official proof status

Area	Estimate	Interpretation
Z0-A design	≈ 90-95%	Current conceptual architecture is ready to close; future changes should be evidence-driven ADRs.
Shared reference runtime	≈ 35-45%	Real components exist; common packaging, recovery and cross-society APIs remain.
Company / Secretary	≈ 45-55% executable reference	Most mature phenotype, but not production deployment.
Trading architecture	≈ 85-90%	Design decisions are relatively mature; code path is only beginning.
Trading runtime	≈ 0-5%	Paper implementation is the next major build.
	≈ 0-20%	

Area	Estimate	Interpretation
Two-society generality proof		Hypothesis and test strategy exist; shared PASS evidence does not yet.
Production	NON-CLAIM	Production is evidence-gated, not a progress-bar percentage.

The strongest defensible summary: World 8 is attempting to turn AI models into replaceable cognitive resources inside persistent, governed digital entities whose identity, authority, truth, effects and development remain accountable outside the model. The decisive next experiment is whether Company and Trading can obey the same laws without forking the Core.

Source and verification map

Project-derived statements in this dossier are grounded in the archived World 6, World 7 and World 8 documents plus recorded repository evidence. External comparison is intentionally limited to official product documentation, standards work and named research papers; it is not an exhaustive novelty search.

Project / archival sources

[P-W6] Entity Ark v6.0 Operational Standard — Implementation Profile 3.1.1; documented release date 12 Aug 2026
Local archival source: [Operation_V06.pdf](#)

[P-W62] World v6.2 public release lineage / DOI
<https://doi.org/10.5281/zenodo.22040348>

[P-W7] World 7: The Genome of a Living Digital World — Technical White Paper v0.2; dated 22 Aug 2026
<https://zenodo.org/records/22058813>

[P-W8] World 8 / Z0-A Complete Architecture — FINAL DESIGN BASELINE / NOT PRODUCTION; frozen 24 Aug 2026
<https://zenodo.org/records/22085394>

[P-SWH] World 8 frozen Software Heritage release
<https://archive.softwareheritage.org/swh:1:rel:5fcdf122c067c17d3763d4cc8d738bff97063bac/>

[P-GH] World 8 public GitHub archival tree
<https://github.com/saeedfaai/World-v6-public/tree/main/world8-z0a>

[GH-SKILL] Commit 632f4cc — versioned proforma Skill executor
<https://github.com/saeedfaai/world-v6/commit/632f4cc0057810b413e7fa55b1b19f66a3a7c06c>

[GH-MASON] Commit d04dd09 — proposal-only Mason v0 runtime
<https://github.com/saeedfaai/world-v6/commit/d04dd09dd4b80f0629bcbb9c65f0355b79d592c9>

[GH-EVAL] Commit 9e057be — independent Evaluator v0 runtime
<https://github.com/saeedfaai/world-v6/commit/9e057bee079843e5e6c144e09e64e08f7b03f52c>

[GH-GATE] Commit 9f6ed7a — Grok MCP / World 8 gateway
<https://github.com/saeedfaai/world-v6/commit/9f6ed7a6a4d62d80e356a4f373aed2d72827f513>

[GH-FARM] Commit aa0e8c9 — recorded multi-provider farm response
<https://github.com/saeedfaai/world-v6/commit/aa0e8c901a4ed7e9d414f9b103a1a08ca6d9eaf7>

External comparison sources

[EXT-OAI] OpenAI, “The next evolution of the Agents SDK”, 15 Apr 2026
<https://openai.com/index/the-next-evolution-of-the-agents-sdk/>

[EXT-LG] LangGraph documentation, Persistence / checkpointing
<https://docs.langchain.com/oss/python/langgraph/persistence>

[EXT-MS] Microsoft Agent Framework, durable extension and checkpoints
<https://learn.microsoft.com/en-us/agent-framework/integrations/azure-functions>

[EXT-TEMP] Temporal platform documentation / durable execution
<https://docs.temporal.io/>

[EXT-ANT] Anthropic, Demystifying evals for AI agents, 9 Jan 2026
<https://www.anthropic.com/engineering/demystifying-evals-for-ai-agents>

[EXT-NIST] NIST NCCoE, Software and AI Agent Identity and Authorization concept paper, Feb 2026
<https://www.nccoe.nist.gov/sites/default/files/2026-02/accelerating-the-adoption-of-software-and-ai-agent-identity-and-authorization-concept-paper.pdf>

[EXT-IETF] IETF Internet-Draft, Architecture for Non-Human Entities, Aug 2026
<https://www.ietf.org/ietf-ftp/internet-drafts/draft-ruvalcaba-nhe-arch-00.html>

[EXT-AD] AgentDevel: Reframing Self-Evolving LLM Agents as Release Engineering, arXiv:2601.04620
<https://arxiv.org/abs/2601.04620>

[EXT-SI] Self-Improvements in Modern Agentic Systems: A Survey, arXiv:2607.13104
<https://arxiv.org/abs/2607.13104>

Prepared as an evidence-bounded technical dossier. No claim of complete prior-art coverage, production readiness, investment return, trading profitability, AGI or autonomous evolution is made. Academic contact: Saeed.farokhi@ut.ac.ir.