

VEHICLE SYSTEMS LAB

SCIENTIFIC CLOSURE AND ARCHITECTURAL RESTART

MADRE-H360 / H360 v1 to H360 v2

Public English Research Edition - Version 1.0 - July 2026

**“We will not protect a version because of the time
invested.**

We will protect the truth of the project.”

H360 v1 is preserved as a learning laboratory.

H360 v2 begins as a life-course architecture.

Prepared under the responsibility of the Director of VEHICLE Systems Lab

DOCUMENT CONTROL

Status and scope of this report

Asset / Decision	Status	Scientific meaning
H360 Legacy Research Population v1	PRESERVED / READ-ONLY	Retained as an experimental population and evidence of population-scale engineering.
H360 v1 independent-layer architecture	CLOSED	No further thematic injection will be treated as the route to complete human life.
Economy T0 derived layer	PRESERVED / NON-FINAL	Kept as a historical derivation; not accepted as a final causal life-course result.
Education lifecycle injection	RESEARCH HOLD	Paused before production because education cannot be assigned independently of biography.
H360 Life-Course Population v2	DESIGN AUTHORIZED	New clean architecture beginning from the first synthetic human.
MADRE query layer	GOVERNED READER	MADRE reads validated layers; it does not assign lives or silently modify the population.

This report is a scientific closure of one generative order, not a declaration that the entire research program failed. H360 v1 demonstrated real capabilities in population generation, batching, storage, lineage, manifests, audit and governed access. It also exposed a structural limitation: human domains cannot be created as independent layers and reconciled afterward without accumulating causal and temporal contradictions.

Evidence rule

All claims in this report are limited to internal laboratory artifacts, counts, audits and experiments. Targets, hypotheses and future objectives are not presented as completed results.

EXECUTIVE SUMMARY

A disciplined reset after a successful scale experiment

VEHICLE Systems Lab reached a major operational milestone: an experimental H360 population consisting of 1,000 batches of 100,000 synthetic records, for a total of 100,000,000. The originals were preserved while derived habitat, memory and economy work was conducted through separate layers. This established practical capability in large-scale generation, organization, storage, verification and governed access.

The decision to stop and restart emerged when the laboratory began designing a complete education lifecycle. Education could not be assigned responsibly from a few demographic or economic attributes. A credible trajectory requires age, cohort, family, territory, infrastructure, language, health, functional capacities, psychological state, work, migration, historical period, institutions, opportunities and interruptions. More importantly, education also changes many of those variables over time.

The same problem applies beyond education. Economy influences schooling, but schooling influences occupation and adult income. Health affects work and education, while work and education affect health. Family creates constraints and support, while economic and educational transitions reshape family life. The architecture therefore needs a shared temporal source of truth rather than a chain of independent injections.

Scientific decision

H360 v1 will remain preserved as a governed learning laboratory and as a large synthetic-data population for bounded pre-investment testing. H360 v2 will begin from the first synthetic human and will scale only after life-course coherence is validated.

1 / ORIGINAL RESEARCH OBJECTIVE

The ambition that remains valid

The core ambition of MADRE-H360 is to create a governed synthetic population that can support auditable research, AI evaluation, multi-perspective reasoning and controlled experimentation without depending on private real-person data. The intended asset is not a set of decorative personas. It is a population whose members possess traceable identities, histories, constraints, relationships, memories and evolving states inside a finite synthetic world.

MADRE is the governed access and memory layer. VEHICLE/H360 is the operator that designs and generates the world and its population. This separation remains fundamental: the query layer must not become the hidden author of the lives it later interprets.

What changed

The product thesis became more precise. The valuable future asset is not “100 million records.” It is “100 million internally coherent, reproducible and queryable life courses.” The v1 work established that the laboratory can reach population scale. The v2 work must establish that the population deserves to be scaled.

What did not change

- The commitment to a fully fictitious governed synthetic population.
- The requirement for immutable originals, versioning, hashes, manifests, ledgers and rollback.
- The goal of population-scale AI evaluation and auditable multi-perspective research.
- The separation between generation, validation, memory access and public presentation.
- The intention to use accelerated computing only after controlled baselines and measurable workloads exist.

2 / VERIFIED OPERATIONAL ACHIEVEMENTS

What H360 v1 proved in practice

- Generation and preservation of a 100,000,000-record experimental synthetic population.
- Organization into 1,000 controlled batches of 100,000 records each.
- A completed inventory audit reporting 1,000 expected batches, zero missing and zero unexpected.
- Profile analysis over 20,000 records distributed across ten batches, with zero invalid JSON in that sample.
- Structured identifiers, public IDs, lineage, manifests, ledgers and SHA-256 verification workflows.
- Private cloud storage and governed metadata-only query routing.
- A planetary habitat research layer and derived economic work without overwriting the original population.
- Experience designing memory, query and API boundaries for MADRE.
- A practical understanding of how fast scale can amplify small modeling errors.

Why these achievements matter

A scientific reset does not erase engineering evidence. The laboratory has already crossed several barriers that often remain conceptual: deterministic generation at large scale, batch control, cloud organization, audit artifacts and preservation of originals. These capabilities are reusable in H360 v2 and reduce execution risk.

Preservation principle

No legacy object should be deleted or overwritten before inventory reconciliation, hash verification and restoration testing are complete.

What H360 v1 did not demonstrate

- Complete or mutually coherent biographies for all 100 million records.
- A plausible joint distribution of family, age, education, occupation, economy, health and territory.
- Institutional capacity in the habitat sufficient to support every assigned trajectory.
- Clinical or psychometric validity of psychological and cognitive proxies.
- Causal compatibility between the economy T0 layer and education, occupation and family history.
- Longitudinal visual identity and contextual consistency across life stages.
- External validity through comparison with representative microdata and independent expert review.
- GPU throughput, latency or cost targets proposed for future NVIDIA experiments.
- Correct population-scale reasoning by MADRE across millions of perspectives.
- Consciousness, sentience or equivalence to real human beings.

These are not minor disclaimers. They define the line between an operational population-scale data experiment and a validated synthetic life architecture. Publishing that boundary increases, rather than decreases, the credibility of the program.

4 / ROOT-CAUSE ANALYSIS

The limitation was not quantity. It was generative order.

The sequence “create identity, inject habitat, inject economy, inject education, add health, add history” can produce contradictions even when each separate layer appears plausible. Human variables are not independent labels. They are coupled processes that cause and modify one another over time.

Example: education

A credible education lifecycle depends on age, cohort, family resources, distance, language, disability, migration, health, digital access, institutional capacity, policies, work obligations, parenthood and crises. Education then changes literacy, occupation, income, networks, expectations and future decisions. A one-way assignment from static attributes is therefore structurally incomplete.

Circular dependencies

- Family economy affects educational access; education affects adult economy.
- Health affects work and schooling; work and schooling affect health.
- Territory affects institutions and opportunity; migration changes territory and opportunity.
- Psychological state affects decisions; lived events and educational experiences modify psychological state.
- Occupation depends on training and markets; occupation changes income, identity, location and family life.

Architectural conclusion

Thematic layers may remain as query and audit views, but they must be derived from one shared temporal life-course source.

5 / ASSET CLASSIFICATION

What is preserved, recalibrated or retired

Asset / Decision	Status	Scientific meaning
Population factory concepts	REUSE WITH REVISION	Generation, batch and deterministic-seed techniques remain valuable.
100M v1 records	PRESERVE	Use as experimental synthetic data, lineage evidence and bounded infrastructure test material.
Public IDs and batch mapping	PRESERVE	Maintain historical traceability; v2 may use a new namespace with explicit crosswalks.
Habitat structure	AUDIT / EXTEND	Retain territorial assets but add capacity, time, institutions and scarcity.
Economy T0	PRESERVE / RE-DERIVE	Keep as history; produce a joint economy-education-occupation v2 result.
Independent education injection	RETIRE	Education becomes a temporal projection of the Life-Course Kernel.
MADRE read API	REUSE WITH GOVERNANCE	Keep read-only separation and extend only after v2 validation.
Mathematical microdynamics	RETAIN WITH SCOPE	Useful for constrained iterative dynamics; not proof of human validity.
Visual prompts	REBUILD	Replace free prompts with structured scene contracts grounded in biography.

Data-conscious synthetic humans: Test before investing

H360 v1 retains a legitimate use inside the VEHICLE Systems Lab ecosystem: a governed, read-only population for data-conscious experimentation before an organization commits expensive compute, storage, integration or product capital.

In this phrase, “data-conscious” means explicit attention to provenance, lineage, uncertainty, access control, audit and limitations. It does not mean consciousness or sentience.

Bounded use cases

- Stress testing ingestion, transformation, validation and recovery pipelines.
- Sizing storage, compression, partitioning, indexing and retrieval systems.
- Testing batch orchestration, manifests, ledgers and hash reconciliation.
- Rehearsing governed metadata access without private real-person records.
- Evaluating dashboards, search, sampling and aggregate interfaces.
- Benchmarking clustering, embeddings, routing and controlled retrieval.
- Detecting contradictions in data models before they become infrastructure costs.
- Preparing synthetic scenarios before purchasing or scaling compute resources.

Public claim boundary

H360 v1 may support data and infrastructure experiments. It must not be marketed as 100 million validated complete human lives.

The new population will be built as life, not as accumulated attributes

World-first principle

Before population generation, H360 v2 must define territories, historical periods, resources, laws, institutions, capacities, languages, services, prices, mobility and scarcity. The world must be able to support the people it contains.

Life-course principle

Each current state must be derived from ordered events: birth, care, schooling, illness, migration, work, relationships, crises, opportunities, learning, aging and memory. Events are the source of truth, not decorative narrative.

Joint-generation principle

Education, occupation and adult economy must be generated jointly. Family, health, psychology and capacities must also interact through time. Independent assignment is prohibited wherever a known dependency exists.

Explicit-uncertainty principle

Every important variable must distinguish base facts, source-derived parameters, model inference, H360-world assumptions, stochastic variation and uncertainty. Synthetic precision must never exceed evidentiary precision.

No automatic diagnosis

Traits, symptoms, behavior and diagnoses must remain separate. Emotional profiles do not authorize mental-health diagnoses, and cognitive proxies do not equal psychometric tests.

Visual-consistency principle

Image, video and 3D systems must receive a structured scene contract derived from biography, age, body, clothing, activity, economy, architecture, climate, relationships and culture.

A temporal life-course kernel with graph structure

For each synthetic person i at time t , the global state may be represented as:

$$X_i(t) = [B_i(t), F_i(t), T_i(t), H_i(t), P_i(t), E_i(t), W_i(t), R_i(t), M_i(t)]$$

B is base biography; F family and household; T territory; H health and function; P psychological state without overdiagnosis; E education; W work and economy; R relationships; M memory.

State transitions are governed by events, institutions, constraints and controlled variation:

$$X_i(t+1) = F_{\theta}(X_i(t), G(t), I(t), A_i(t), \epsilon_i(t))$$

$G(t)$ is the social and territorial graph, $I(t)$ represents institutions and their capacity, $A_i(t)$ contains actions and events, and $\epsilon_i(t)$ is controlled individual variation. Current snapshots are derived; event histories remain auditable.

Storage and computation model

- Immutable event ledger as the primary longitudinal source.
- Derived snapshots for efficient current-state queries.
- Multiplex graph for family, household, territory, education, work and social relations.
- Capacity ledgers for schools, hospitals, housing, transport and employment.
- Versioned parameter registry with source, transformation, uncertainty and limitations.
- Deterministic random streams partitioned by person, household, institution and event domain.
- Constraint engine capable of halting generation when contradictions exceed thresholds.

The first H360 v2 human must be explainable before being scalable

The new Human 1 will not be approved because the record contains many fields. Approval requires an internally coherent, reproducible and visualizable life.

- Identity, age and cohort compatible with the calendar and history of the world.
- Parents, caregivers, siblings and household with valid reciprocal relationships.
- Territory, housing, services and institutions that actually exist and have capacity.
- Complete non-decreasing chronology of relevant events.
- Education compatible with access, language, family, health, economy and institutional capacity.
- Work and economy compatible with age, competencies, labor market and educational history.
- Health and disability modeled without contradiction or automatic stigma.
- Memories derived from lived events, including controlled imperfection and forgetting.
- Psychological state coherent with history without unsupported diagnosis.
- Social relationships with reciprocity and temporal validity.
- Scene contract capable of producing consistent visual representations across ages.
- Deterministic re-execution and byte-equivalent canonical output under the same version and seed.
- Complete human review and uncertainty report.

Scale becomes an experiment, not a celebration

1. One synthetic human - full life-course audit.
2. One reciprocal household - family and resource coherence.
3. One community - institutions, services, capacity and social graph.
4. 1,000 humans - distribution, edge cases and reproducibility.
5. 10,000 humans - cohort, household and territorial calibration.
6. 100,000 humans - city-scale capacity and performance.
7. 1 million humans - distributed generation and query baseline.
8. 10 million humans - infrastructure, graph and statistical stress test.
9. 100 million humans - only after all prior gates and independent review pass.

Mandatory stop conditions

- Age-stage or chronology contradictions above the approved threshold.
- Broken reciprocal family or household relations.
- Institutional capacity exceeded without modeled alternatives.
- Distribution drift outside source-calibrated uncertainty bands.
- Bias patterns caused by unsupported penalties or stereotypes.
- Non-deterministic canonical outputs under identical version and seed.
- Untraceable parameters, missing provenance or unresolved hash mismatch.
- Visual outputs that invent incompatible identity, context or cultural signals.

The revised problem is a stronger accelerated-computing problem

The v2 challenge is not merely to store or scan 100 million rows. It is to generate, validate, update and query a large temporal graph under deterministic lineage, capacity limits, uncertainty and cross-domain constraints. That creates measurable workloads suitable for controlled accelerated-computing experiments.

Candidate workloads

- Parallel event processing across persons, households, institutions and territories.
- Constraint validation for age, chronology, capacity, reciprocity and resource compatibility.
- Graph-neighborhood retrieval and temporal aggregation.
- Embedding, clustering and governed multi-perspective sampling.
- Batch visual-consistency evaluation from structured scene contracts.
- GPU-versus-CPU benchmarks for throughput, latency, energy and cost per generated life-year.
- Triton-based serving experiments for validated aggregate queries.
- NeMo or related tooling only where the experiment has explicit metrics and reproducible baselines.

Non-endorsement statement

This research report and public page are prepared for possible NVIDIA Innovation Lab consideration. They do not claim NVIDIA endorsement, funding, technical validation or partnership.

Legacy preservation protocol

10. Create a master inventory of all local and cloud objects associated with H360 v1.
11. Generate SHA-256 hashes and counts by prefix, batch, manifest, ledger, source code and document.
12. Register versions, dates, dependencies, validation state and responsible owner.
13. Apply technical and documentary read-only controls to legacy datasets.
14. Separate source code, generated artifacts, results, secrets, environments and temporary data.
15. Maintain at least two verified copies and one independent cold backup.
16. Create an evidence index that allows historical review without reactivating every service.
17. Prohibit deletion until inventory reconciliation, restoration testing and director approval are complete.

Required architecture decisions

- ADR-H360-000: closure of the independent thematic-injection architecture.
- ADR-H360-001: authorization of the Life-Course Population v2 architecture.
- ADR-H360-002: H360 v1 retained for bounded data-conscious experimentation.
- ADR-H360-003: MADRE remains a governed reader after validation and version freeze.

Truthful positioning protects the project

The reset prevents a common failure in synthetic-data programs: continuing to add fields until a record looks human while its history remains impossible. H360 v2 must not use complexity as decoration. Every added domain must improve explanatory, causal, temporal or operational coherence.

Public communication rules

- Distinguish verified achievements from targets and proposals.
- Describe v1 as an experimental synthetic research population, not a complete civilization model.
- Explain that data-conscious does not mean sentient.
- Publish limitations alongside capabilities.
- Do not expose raw population records, seeds, private factories or restricted prompts.
- Do not promise GPU results before benchmark execution.
- Do not use real-person identity or private data to increase apparent realism.

Commercial discipline

The valid near-term offer is pre-investment testing of data architecture, governance and synthetic scenarios. The long-term offer is a validated life-course population for AI evaluation and research. These must remain separate until evidence connects them.

Director-authorized scientific decision

VEHICLE Systems Lab closes H360 v1 as the active route to a complete human-life architecture. The population, code, reports, APIs, ledgers and derived layers are preserved as laboratory evidence and reusable engineering assets. They are not deleted, overwritten or retroactively presented as something they did not prove.

The laboratory authorizes H360 v2 as a clean reconstruction beginning from Human 1, household 1 and world constraints. Scale will be earned through validation gates rather than assumed from prior volume.

Closing declaration

We will not protect a version because of the time invested. We will protect the truth of the project. H360 v1 is preserved as a learning laboratory. H360 v2 begins as a life-course architecture, carrying forward every objective, omission, constraint and lesson discovered through H360 v1.

This decision accepts the cost of rebuilding because future correction across 100 million interconnected lives would be more expensive, less transparent and less scientifically defensible than rebuilding from a correct generative foundation.

Internal evidence classes supporting this closure

- Population batch and manifest inventories.
- Source ZIP inventories and upload verification records.
- Profile analyses over controlled record samples.
- Mitochondrial-memory and query-layer manifests.
- Cloud storage structure and private API deployment records.
- Economy T0 generation, audit and ledger artifacts.
- Education-lifecycle research brief and paused bootstrap architecture.
- Mathematical preprints and numerical experiment reports.
- Public NVIDIA-oriented research justification and architecture notes.
- Scientific closure report, architecture decision records and future v2 specifications.

A future independent review should receive a frozen evidence bundle with hashes, source locations, creation dates, responsible operator and a statement identifying which claims each artifact supports.

Appendix B - Terminology

Term	Definition
H360 v1	The preserved experimental 100-million-record population and its legacy derived research layers.
H360 v2	The authorized life-course architecture that will restart from the first synthetic human.
Life-Course Kernel	The temporal generative system from which current human domains are derived.
Derived layer	A versioned output linked to immutable source identity and events without overwriting originals.
Data-conscious	Explicitly designed around provenance, lineage, uncertainty, governance and data limitations; not sentence.
MADRE	The governed reader and integrator of validated population layers, not the author of individual lives.
Scene contract	A structured visual representation request grounded in biography, context and uncertainty.

FINAL RECORD

Scientific closure record

Asset / Decision	Status	Scientific meaning
Decision	APPROVED	Close active development of the independent-layer H360 v1 architecture.
Preservation	MANDATORY	Keep v1 population and evidence read-only with verified hashes and backups.
Current valid use	AUTHORIZED	Data-conscious synthetic-human data and infrastructure testing before investment.
New architecture	AUTHORIZED	Begin H360 v2 Life-Course Architecture from Human 1.
Scale to 100M	NOT AUTHORIZED YET	Requires successful scientific, temporal, statistical, technical, ethical and visual gates.

Director: Roberto David Borda Milan

Organization: VEHICLE Systems Lab

Project: VEHICLE / MADRE / H360

Record date: July 2026

Document version: 1.0 - Public English Research Edition