

SYSTEMICA ATLAS V2.1.2

Executive Digest

Public Release 2.1.2 - 21 July 2026

Joseph Maxwell

Founder-originated method; AI-assisted implementation; human-controlled authority

FOUNDER-AUTHORISED PUBLIC RELEASE - CLAIM BOUNDARIES REMAIN IN FORCE

Contents

The proposition	2
The founder and the method	3
Where generative AI is - and is not - involved	4
Drift, state and recovery	5
What has been demonstrated	6
Measured telemetry	7
Materials evidence	8
Three engineering missions	9
What the customer receives	10
Commercial route and boundary	11

The proposition

Systemica helps engineers decide what modelling evidence is trustworthy enough to justify the next expensive action.

GIL-EDK is the initial commercial wedge; state integrity, CTAP and governed recovery are the defensible expansion layer.

Joseph Maxwell developed the underlying method: map an unclear system, expose contradictions, test it, decide what deserves trust, and determine the next action. Systemica formalises that method in software, evidence, memory and authority controls.

The founder and the method

From model optimisation to governed evidence

Historical artefacts explain the problem transition; accepted runtime evidence begins at GIL-EDK.



Figure 1: Method-to-product lineage.

Joseph Maxwell's route began in materials engineering and structured experimentation. An early neural-network workflow systematically searched 127 descriptor combinations for shear-modulus prediction. The apparent winner moved when the process was repeated, while code inspection exposed target scaling, test-set reuse and physically adjacent descriptors that an impressive score did not disclose. GIL-EDK grew from that gap: it audits which relationships survive controlled stress and recurrence. Systemica governs what happens next - what may be promoted, what must be held, and how the last accepted state can be recovered.

The same state-space logic later became a wider research programme around constraint, navigability, drift and recovery; it remains conceptual context rather than product evidence. GPT/ChatGPT accelerated critique; Codex performed substantial implementation and testing. Final technical and publication authority remains human.

Where generative AI is - and is not - involved

Where generative AI is - and is not - involved

Core execution remains separated from optional assistance and human authority.



Measured evidence cannot be rewritten or promoted by a generative route.

Figure 2: Generative-AI authority boundary.

The tested core is generative-AI-detached, not universally AI-free. GIL uses conventional fitted ML. Project Intelligence is an optional generative route. No generative route can promote evidence or accepted memory.

Drift, state and recovery

Drift, state and governed recovery

A declared change is observed, validated and either continued, held or recovered without allowing the model to certify itself.



Controlled evidence: 16 cases; 8 HOLD; 8/8 held writes refused; 0 live LLM runs.

Figure 3: Drift, state and governed recovery route.

Systemica tests whether a conclusion deserves trust now, then prevents that accepted state from silently degrading later. In 16 controlled cases, 8 threshold breaches produced **HOLD**, all 8 held writes were refused, and rollback recovered the accepted fingerprint. These were bounded integration fixtures with **0 live LLM runs**, not enterprise drift-detection evidence.


What has been demonstrated

- Materials Project audits to 10,982 rows across shear and bulk;
- independent EDK kernel with two-module parity;
- typed state, transition and promoted-memory bridge;
- three cross-domain Design-3 missions ending in bounded HOLD;
- credential-free offline acceptance of representative core routes.

None of this is external scientific validation or product readiness.

Measured telemetry

Measured local compute and synthetic memory scale




Measured on one local machine. HMA scale is synthetic fixture evidence. No human-time, cost-saving or production-throughput claim.

Figure 4: Compute and memory telemetry.

The locked 2,000-row GIL route completed in **192.83 seconds** with **470 MiB** sampled peak process-tree RSS. Human-time and cost savings remain unmeasured.

Materials evidence



All eight scale points reported zero stable descriptors. Source: accepted MP expansion run; limitation: local computational evidence, not external validation.

Figure 5: Materials Project scale evidence.

Prediction remained strong at scale, while no descriptor earned stable status. That distinction between predictive signal and earned stability is the current GIL-EDK product lead.

Three engineering missions

Three real missions, three bounded HOLD outcomes

The workflow transferred across thermal interface, fluid topology and porous thermal-management problems.



Figure 6: Three bounded engineering missions.

Furnace, dual tank and heat shield each produced traceable candidates, exposed weaknesses and narrowed the next physical question without unsupported promotion.

What the customer receives

1. A bounded problem contract.
2. A controlled run with declared assumptions.
3. A weakness and contradiction register.
4. Traceable candidate and route lineage.
5. A **HOLD**, **REVISE** or **CONTINUE** decision packet.
6. The smallest justified next test.

This can help a team make an expensive next step more inspectable. Measured customer benefit, willingness to pay and workflow savings remain future evidence gates.

Commercial route and boundary

Commercial adoption ladder

Service-led entry builds evidence; repeatability and licensing are later gates.




Figure 7: Bounded adoption route.

The next gates are one independent-source replication, one bounded external pilot and one prospectively measured workflow comparator. This release is founder-authorised; external legal/IP review and independent validation remain absent.

State: SYSTEMICA_ATLAS_V2_1_2_FOUNDER_AUTHORISED_PUBLIC_RELEASE_ACCEPTED