Independence as an Architectural Property: A Research Direction for Multi-Source Cryptographic Entropy

Why Future Entropy Architectures Should Design for Verifiable Spatial Independence Rather Than Assume It

By Jorge Enrique Flores Montano

Founder, JM Automated Solutions

~MILO™ — Modular Intelligent Learning Orchestrator (patent pending)

May 2026

ORCID iD: 0009-0003-1859-8418

DOI: 10.5281/zenodo.20117647

Public reference: https://github.com/jmontano1/milo-architecture

Abstract

True random number generators that seed cryptographic systems are evaluated on two properties: the entropy density of the physical noise being sampled and the *independence* of the samples produced. NIST SP 800-90B specifies that a multi-source entropy generator may combine multiple noise sources only when the sources are *independent* — that is, when the joint distribution of source outputs is the product of the marginals. In current practice, independence is achieved by substrate engineering: chaotic laser arrays and microcomb-based parallel chaos sources are *designed* to be independent at the optical substrate, and the independence claim is supported by inter-channel cross-correlation measurements at the device level. This paper argues that the next architectural direction for multi-source entropy generation in critical-infrastructure cryptography is to *design for verifiable independence at the joint-observation level* — to build entropy architectures in which the independence of sampled noise is a measurable architectural property rather than an assumed substrate property. The position is framed as a research direction, not a specific implementation: this paper argues*that* the architectural problem is worth pursuing, not *which* apparatus solves it. The argument is positioned within the NIST SP 800-90B framework, distinguished from existing single-source thermal, single-camera optical-chaos, and parallel-optical TRNG architectures, and connected at the architectural level to MILO's sensor-receptor subsystem.

Keywords: cryptographic entropy, true random number generation, NIST SP 800-90B, multi-source entropy, independence verification, architectural design.

Highlights.

- Proposes*independence-verifiable joint extraction* as a measurable architectural property of multi-source cryptographic entropy generators, distinct from substrate-engineered approaches.
- Positions the research direction within the NIST SP 800-90B framework as a design-time constraint rather than a post-deployment validation step.

- Distinguishes the proposed architectural class from single-source thermal TRNGs, single--camera optical-chaos sources, and parallel-optical TRNG architectures via two falsification criteria.
- Connects the architectural class at the public-architecture level to MILO's sensor-receptor subsystem; specific apparatus is not proposed.

Index Terms: cryptographic entropy, true random number generation, NIST SP 800-90B, multi-source entropy, independence verification, joint-observation extraction, signal processing, critical-infrastructure cryptography.

Plain Language Summary. Cryptographic systems that protect critical infrastructure — power grids, secure communications, financial transactions, defense command-and-control — depend on a continuous supply of high-quality random numbers seeded by physical noise sources. When multiple noise sources are combined, the resulting random numbers are only as strong as the *independence* between sources. Today, independence is achieved by building sources that are independent at the substrate level (for example, separate laser chaos channels). This paper argues that the next architectural step is to design entropy systems where independence is not assumed but *measured* at the joint-observation level — built in, verifiable, and auditable. The argument is a research-direction position paper, not a claim about a specific apparatus.

Relevance to U.S. National Interest. Cryptographic entropy sourcing underpins the trustworthiness of every protected communication and authentication system in U.S. critical infrastructure. The U.S. Department of Energy's Genesis Mission identifies AI-enabled advancements in critical-infrastructure security as a national-importance priority; the architectural direction proposed here addresses a foundational layer of that priority.

Status of claims. This is a position paper proposing an architectural*research direction*; it does not present a specific apparatus or empirical validation. The NIST SP 800-90B [1] independence requirement, the single-source thermal TRNG architectures[2], [3], the single-camera optical-chaos source [4], the parallel-optical TRNG architectures[5], [6], and the wearable-sensor entropy extraction work [7] are established external references and are described as such. The architectural class advanced in §3 — independence-verifiable joint extraction at the multi-source observation level — is the author's original proposal, framed as a falsifiable research direction (see §3, "Falsification criteria"). Compatibility with MILO's sensor-receptor subsystem[8] is illustrative; an apparatus-level instantiation and empirical validation are forthcoming work and are not claimed here. This manuscript is a preprint prior to peer review.

About MILO. MILO (Modular Intelligent Learning Orchestrator) is a patent-pending adaptive AI orchestration architecture for high-consequence critical-infrastructure environments. The full architectural reference — eight structural principles, eight operational integrity constraints, the unifying viability principle, and the trademark/patent status — is maintained as a single canonical document at https://github.com/jmontano1/milo-architecture (concept DOI: 10.5281/zenodo.20117025). Author contact: jmontano@jmautomated.com.

1. Introduction

A cryptographic random number generator is no stronger than the entropy source that seeds it. The quality of the entropy source is bounded by two properties: the *entropy density* of the physical phenomenon being observed, and the *independence* of the noise samples produced by the observation. NIST SP 800-90B [1] specifies that multi-source entropy generators may combine multiple noise sources only when the sources are independent — that is, when the joint distribution of the source outputs is the product of the marginals. The independence requirement is operationally non-trivial; in practice, most multi-source entropy generators have either (a) used multiple observations of the same physical phenomenon (e.g., multiple readings of resistor thermal noise), (b) used multiple physically distinct phenomena (e.g., resistor noise mixed with CMOS sensor noise mixed with avalanche-diode noise), or (c) engineered the physical substrate so that channels are independent by construction (e.g., chaotic laser arrays).

The thermal-source TRNG literature has explored three principal architectures. The first samples Johnson noise in electronic components — typically resistors observed through a programmable analog-to-digital converter; the canonical recent example is Matsuoka's 2021 PSoC TRNG [2], in which four variant designs pass NIST SP 800-22 statistical testing. The second samples thermal noise in CMOS image sensor pixels — Vault12's TrueEntropy[3] is a commercial example using the standard smartphone CMOS camera to extract thermal pixel noise. The third samples optical chaos in a controlled scene — Cloudflare's LavaRand [4] uses a wall of approximately one hundred lava lamps observed by a single camera, with related Cloudflare deployments using wave motion (Lisbon office) and hanging mobiles (Austin office). Each architecture is single-source or single-observation; multi-device deployment produces parallel instances of a single-source pattern rather than a multi-source generator in the SP 800-90B sense.

A separate literature has developed parallel and multi-channel TRNG architectures in the optical / laser regime. Recent work has demonstrated massively parallel chaos based on chaotic microcombs [5], achieving an aggregation rate of 3.84 Tbps with inter-channel correlation below 0.04 across thirty-two channels. Two-dimensional chaotic laser arrays have achieved 4×4 independent channels with all cross-correlation coefficients below 0.05 [6]. The parallel-optical literature operates within a specific solution pattern:*independence is engineered into the optical substrate itself*, and the independence claim is supported by device-level cross-correlation measurements after the fact.

This paper takes a step back from the question of *which apparatus to build* and asks a different question:*what should the architectural property of independence look like in a multi-source entropy generator*? The position developed here is that independence should be a *designed-for and measurable property of the joint observation*, not an assumed property of the substrate. The argument is framed at the architectural-principle level; specific apparatus, sensors, geometries, sampling configurations, and extraction algorithms are not proposed in this paper. The intent is to motivate a research direction, not to claim a particular instantiation.

2. Background and Related Work

The prior art across thermal, optical, and parallel TRNG architectures can be organized along two axes: how many sources are sampled, and whether independence is assumed or measured.

Single-source thermal noise from electronic components. Matsuoka 2021 [2] presents four PSoC TRNG variants, three of which use the internal resistors of a programmable gain amplifier as the noise source through a delta-sigma analog-to-digital converter. All four pass NIST SP 800-22 [10]. The architecture is single-source per device; the entropy density is bounded by the thermal noise variance of one resistor under one observation. Multi-device deployment produces multiple instances of the same architecture, not a multi-source entropy generator in the SP 800-90B sense.

Single-camera thermal-pixel noise from CMOS sensors. Vault12's TrueEntropy[3] uses the thermal noise of CMOS image sensor pixels in a standard smartphone camera; the application has been evaluated under standard NIST randomness testing and is openly available as iOS source code. The architecture is single-camera; the entropy density is bounded by the variance of one CMOS sensor's pixel-noise distribution.

Single-camera optical-chaos observation. Cloudflare's LavaRand [4] observes approximately one hundred lava lamps from a single camera and mixes the resulting pixel-stream with the production data center's local entropy. The framework's own technical disclosure notes that LavaRand is a secondary independent source — the primary is RDRAND — and that the camera-mixed entropy is one input among several. The architecture is single-camera; multi-deployment instances (Lisbon, Austin) use different physical phenomena under the same single-camera pattern.

Wearable sensor entropy extraction. Recent work [7] has explored wearable-device sensors (heart rate variability, accelerometer/gyroscope, photoplethysmography, electrodermal activity, skin temperature) as entropy sources for cryptographic key generation. The architecture varies per device but operates within the single-wearable, single-physical-modality paradigm; multi-device deployment is parallel-single-source rather than multi-source-spatial.

Parallel optical/laser TRNG. The optical chaos literature has developed parallel TRNG architectures using chaotic microcombs[5] (3.84 Tbps, inter-channel correlation< 0.04), chaotic laser arrays[6] (4×4 channels, cross-correlation< 0.05), and spatial-light-modulator + CMOS-camera spatiotemporal chaos. These architectures operate within a specific approach to independence:*the substrate is engineered so that channels are independent by construction*. Independence is then verified post-hoc through cross-correlation measurements at the device level.

The pattern across the prior art is consistent. Single-source architectures rely on the variance of one physical observation; multi-source architectures in the optical regime rely on substrate-level engineered independence.**No published architecture, to the author's working knowledge of the field, treats independence as a measurable property of the joint multi-source observation itself rather than as a substrate property to be verified after the fact.** This is the architectural gap the present paper identifies as a research direction.

3. Independence as an Architectural Property

The position developed here is that independence in a multi-source entropy generator should be designed for at the architectural layer — built into the system's design specification so that independence can be observed, measured, and audited at the joint-observation level. The argument has three components.

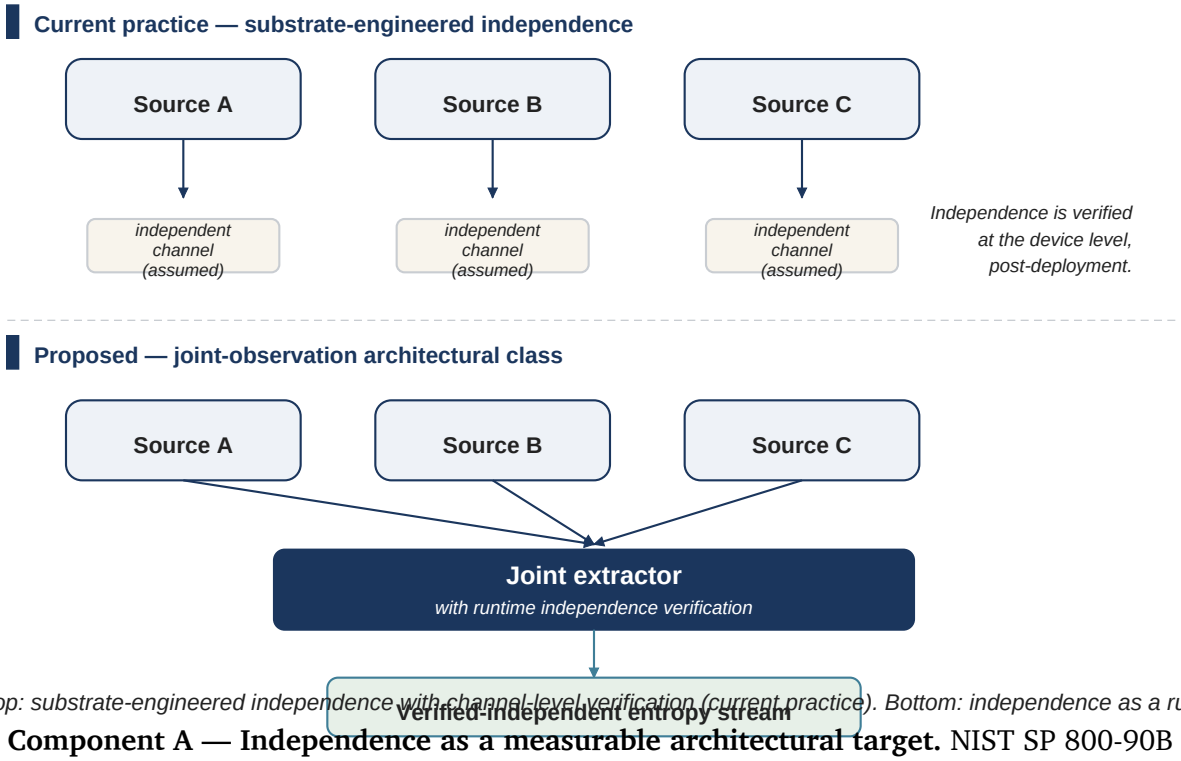

ptographic entropy. Top: substrate-engineered independence with channel-level verification (current practice). Bottom: independence as a runtime architectural pro

Component A — Independence as a measurable architectural target. NIST SP 800-90B requires that multi-source entropy generators combine sources only when the sources are independent. In current practice, this requirement is met by substrate engineering: the system is built such that, by physical construction, the sources are independent. Independence is then *checked* by cross-correlation measurement, but the architecture itself does not *enforce* independence; it assumes it. The architectural alternative is to design the system so that independence is the *property the joint extraction is engineered to verify and preserve*, rather than the property the substrate is engineered to provide. In this alternative architecture, the extraction layer's central task is independence verification, not noise mixing.

Component B — Independence at the joint-observation level, not the per-source level. The current parallel-TRNG pattern verifies independence at the source level — channel X has cross-correlation below threshold C with channel Y. The architectural alternative proposed here is to verify independence at the *joint observation*: given a set of simultaneous observations from multiple sources, the extraction layer separates jointly-observable shared components from per-source residual components, retains the residuals whose joint distribution can be verified independent, and discards the rest. The shift is from *trusting that the substrate produces independence* to *measuring independence in the observation and conditioning the output to preserve only the verifiably independent portion*.

Component C — Conformance with NIST SP 800-90B as a design target, not a test pass. NIST SP 800-90B specifies design principles, health tests, and conditioning functions for entropy sources. Architectures designed under the principle proposed here would treat the SP 800-90B health-test suite as an *architectural conformance target*: the health tests operate on the joint output as an integral part of the architecture, not as an external validation step after deployment. The independence-verifiable property is built into the system rather than tested for after the fact.

The three components together describe an architectural class — multi-source entropy generators in which independence is a designed-for, measurable, and architecturally enforced property of the joint observation. The paper does not claim a specific apparatus in this class; the position is that the class itself is worth pursuing.

Mathematical framing. In information-theoretic terms, the proposed architectural class targets a joint extraction $f(X_1, \ldots, X_n) \to Y$ such that the pairwise mutual information $I(Y_i ; Y_j)$ is bounded above by a configurable verification threshold for all retained output components Y_i, Y_j. The extraction is conditioned on the joint observation $(X_1, \ldots, X_n)$ so that shared components — components present in more than one source — can be estimated and conditioned away before the residual independently-supported components are retained. The independence claim becomes a measurable property of the retained output Y under the architecture's runtime, not an assumed property of the source substrate X. NIST SP 800-90B's health-test suite operates on Y as part of the architecture's continuous operation, not as a one-time certification event. The architectural class is parameterized by the verification threshold, the extraction function family, and the conformance target; specific instances of those parameters define specific architectures within the class.

Falsification criteria. The position developed here is falsifiable on two axes. *(1)* If a multi-source entropy architecture built under the proposed class produces an independent-entropy output rate consistently lower than that of an otherwise-equivalent substrate-engineered architecture in the same physical regime, the operational value of the proposed class is contested. *(2)* If theoretical bounds on extractable independent entropy from a joint observation can be shown to be saturated by current substrate-engineered approaches — that is, if the joint-observation approach offers no measurable architectural benefit beyond what substrate engineering already achieves — the necessity of the proposed class is contested. The architectural class is therefore an empirical research target, not a theoretical claim that is immune to disproof.

Why the field has not already taken this direction. A reasonable objection to the position above is that, if it were valuable, the field would already have moved this way. Three factors plausibly account for the current absence. *First*, in the optical/laser regime where parallel multi-channel TRNG architectures have been most actively developed, the engineered substrate (chaotic microcombs, multimode laser arrays) admits natural decorrelation through device physics; the operational pressure to design verification into the joint extraction has been low because the substrate produces near-independence by construction. *Second*, in the thermal-source regime, single-source architectures (resistor thermal noise, CMOS pixel noise) have been operationally adequate for many cryptographic-engineering applications, and the cost-benefit of moving to multi-source has not been articulated in the entropy-quality terms developed here. *Third*, the independence-verification layer described in §3, Component B requires statistical machinery (joint-distribution estimation, conditional independence testing under SP 800-90B health-test constraints) whose tooling has matured significantly in recent years but was not standard practice in earlier TRNG design epochs. None of these explain why the direction is closed; they explain why the direction has been open and unattended.

4. MILO as the Working Laboratory

The architectural direction described in §3 — independence as a designed-for, measurable, architecturally enforced property — is being explored in practice on MILO (Modular Intelligent Learning Orchestrator), the patent-pending adaptive AI orchestration architecture developed by the author[8]. MILO is the working substrate for the architectural class proposed here. Three architectural properties of MILO map directly onto the requirements of the proposed class. *(i)* A sensor-receptor abstraction admits multiple physical-input sources observing a common phenomenon and emits their samples through an audit-first signal bus. *(ii)* The audit-first command-and-signal substrate persists every command before dispatch and every signal before fanout, ensuring that every entropy-capture event and every quality-metric report is auditable end-to-end. *(iii)* A monitoring subsystem publishes entropy-quality reports backed by persisted events, so the cryptographic consumer's post-hoc verification of *which capture events contributed to which seeds* is a mechanical property of the architecture rather than a logging convention.

MILO is not the only orchestrator that could host the architectural class proposed in §3, but it is the substrate on which the author is developing the direction. The compatibility described above is therefore not hypothetical — it is the working hypothesis under which MILO has been submitted under the U.S. Department of Energy Genesis Mission[9].

5. Implications and Discussion

For *design*, the position implies that future multi-source entropy architectures should target measurable independence at the joint-observation level as a design specification rather than as a post-deployment validation step. The implication is not that current substrate-engineered approaches are unsafe; it is that an architectural class exists in which the independence claim becomes a property of the design rather than a property of the test report.

For *evaluation*, the position implies that NIST SP 800-90B independence verification — operationally non-trivial in any multi-source architecture — could become an integral part of the architecture's runtime rather than a one-time validation event. The verification of independence becomes a continuous property of the joint output, not a snapshot taken at certification time.

For *deployment*, the position implies that high-consequence cryptographic applications — including those in critical infrastructure protection, secure communication for industrial control systems, and entropy seeding for autonomous-system command-and-control — should evaluate the architectural class proposed here as a candidate research direction for their next-generation entropy supply. The DOE Genesis Mission's emphasis on advanced manufacturing, grid reliability, and human-in-the-loop AI in high-consequence domains[9] is a natural deployment context.

5.1 Limitations

This paper does not present an apparatus, an empirical evaluation, or a worked instance of the proposed architectural class. The position is presented at the architectural-principle level only; specific apparatus, sensor configurations, geometries, sampling protocols, and extraction algorithms are not proposed. The compatibility argument with MILO's sensor-receptor subsystem[8] is illustrative, not load-bearing. Two specific empirical questions are unresolved: (i) whether joint-extraction architectures of the proposed class can produce independent-entropy throughput competitive with substrate-engineered parallel-optical TRNGs [5], [6] under comparable optical regimes; and (ii) whether the runtime independence-verification machinery imposes throughput penalties incompatible with cryptographic-seeding cadences in high-consequence deployments. Both are forthcoming work and bound the falsifiability criteria in §3.

6. Conclusion

This paper has proposed an architectural research direction for multi-source cryptographic entropy generation: that independence should be a *designed-for, measurable, and architecturally enforced property* of the joint observation, rather than an assumed property of the substrate. The position is positioned within the NIST SP 800-90B framework for multi-source entropy generators [1] and is distinguished from the current pattern of substrate-engineered parallel chaos sources [5], [6], single-source thermal TRNGs [2], [3], single-camera optical-chaos sources [4], and wearable sensor entropy extraction[7]. The position is presented at the architectural-principle level only; specific apparatus, sensor configurations, geometries, sampling protocols, and extraction algorithms are not proposed in this paper. The architectural connection to MILO's sensor-receptor subsystem[8], submitted under the U.S. Department of Energy's Genesis Mission[9], illustrates compatibility at the public-architecture level. Related architectural directions — latency-aware authentication in industrial control, supervisory primacy for human-in-the-loop AI orchestration, structural principles for adaptive AI architecture, and the discipline of viability under non-stationary conditions — are developed in related work by the author.

Data Availability

All architectural materials, source manuscripts, the reference implementation, and accompanying figures are openly available at https://github.com/jmontano1/milo-architecture and permanently archived at Zenodo (DOI: 10.5281/zenodo.20117025). No private datasets are referenced; the architectural framework itself is the subject of this paper. Patent rights for the underlying MILO software architecture are reserved; the ~MILO trademark is held under USPTO Serial No. 99706004 (intent-to-use, Class 009).

References

[1] M. S. Turan, E. Barker, J. Kelsey, K. A. McKay, M. L. Baish, and M. Boyle, *Recommendation for the Entropy Sources Used for Random Bit Generation*, NIST SP 800-90B, National Institute of Standards and Technology, Jan. 2018. doi:10.6028/NIST.SP.800-90B

[2] M. Matsuoka, "A true random number generator that utilizes thermal noise in a programmable system-on-chip (PSoC)," *International Journal of Circuit Theory and Applications*, vol. 49, no. 10, pp. 3354–3367, May 2021. doi:10.1002/cta.3046

[3] Vault12, *TrueEntropy: High Volume Thermal Entropy Generator for iOS*, [Online]. Available: https://github.com/vault12/TrueEntropy and https://apps.apple.com/us/app/trueentropy/id1299321174

[4] Cloudflare, Inc., "LavaRand in Production: The Nitty-Gritty Technical Details," *Cloudflare Blog*, 2017–2024. [Online]. Available: https://blog.cloudflare.com/lavarand-in-production-the-nitty-gritty-technical-details/

[5] B. Shen, H. Shu, W. Xie, et al., "Harnessing microcomb-based parallel chaos for random number generation and optical decision making," *Nature Communications*, vol. 14, art. 4590, 2023. doi:10.1038/s41467-023-40152-w

[6] Q. Chen, X. Tang, M. Xu, F. Zhang, F. Zhang, Y. Guo, M. Pu, and X. Luo, "Spatiotemporal chaos based on a multimode laser for parallel physical random number generation," *Optics Letters*, vol. 51, no. 8, pp. 2284–2287, 2026. doi:10.1364/OL.589666

[7] M. Svarcmajer, M. Kohler, Z. Krpic, and I. Lukic, "Entropy Extraction from Wearable Sensors for Secure Cryptographic Key Generation in Blockchain and IoT Systems," *Sensors*, vol. 25, no. 17, art. 5298, 2025. doi:10.3390/s25175298

[8] J. E. Flores Montano, *MILO (Modular Intelligent Learning Orchestrator)*, JM Automated Solutions. Patent pending. Submitted under the U.S. Department of Energy Genesis Mission, 2026.

[9] U.S. Department of Energy, "The Genesis Mission: Transforming Science and Energy with AI," Office of the Under Secretary for Science, Executive Order 14363, November 2025. [Online]. Available: https://www.energy.gov/genesis

[10] L. E. Bassham et al., *A Statistical Test Suite for Random and Pseudorandom Number Generators for Cryptographic Applications*, NIST SP 800-22 Rev. 1a, National Institute of Standards and Technology, Apr. 2010. doi:10.6028/NIST.SP.800-22r1a

About the author

Jorge Enrique Flores Montano (ORCID iD: 0009-0003-1859-8418; jmontano@jmautomated.com) is the founder of JM Automated Solutions and the inventor of MILO. A full biography is maintained at https://www.milo-usa.com/jorge-enrique-flores-montano.html.

Conflict of Interest and Funding Disclosure

The author is the inventor of MILO (patent pending) and the founder of JM Automated Solutions. The architectural research direction proposed in this paper is a contribution from a working development program in which the author retains sole authorship and inventive interest. No external funding was received for the preparation of this manuscript.

Appendix A — About MILO

MILO (Modular Intelligent Learning Orchestrator) is a patent-pending adaptive AI orchestration architecture organized into discrete, single-responsibility subsystems under a strict separation-of-concerns discipline. An audit-first command-and-signal substrate persists every command before dispatch and every signal before fanout, producing an append-only audit trail that survives arbitrary process termination. The architecture is designed for *viability* — operational continuity under non-stationary conditions — rather than for prediction accuracy against an expected future.

Eight structural principles

Six are established physical, informational, control-theoretic, and statistical laws applied as architectural design constraints; two are original frameworks proposed by the author for the operator-cognitive performance layer of high-consequence systems.

1. **Second Law of Thermodynamics** — entropy treated as an architectural diagnostic signal, not a fault to be suppressed.
2. **Ashby's Law of Requisite Variety** — a regulator must possess variety at least equal to the system it regulates; implemented as a fleet of specialist agents matching the operational domain.
3. **Shannon Information Theory** — variance reduction occurs at the signal-carrier level, not redundantly at each consumer.
4. **Principle of Least Action — Single-Target Dispatch** — every command has one explicit target; no implicit resolvers, no opaque dispatchers.
5. **Lyapunov-Style Bounded Response** — every adaptive subsystem admits an explicit halt-and-resume pathway; adaptation that drifts unboundedly is failure, not learning.
6. **Power-Law Distribution Architecture** — engineered for the 99th-percentile event, not the median.
7. **Individual-Baseline Variance Modeling** *(original framework)* — operator-layer interventions calibrated against the individual's own established performance baseline, never a population norm. Design-stage; pending empirical validation.
8. **Precision Perturbation Without Variance Compression** *(original framework)* — operator-layer interventions shift probability mass toward high-reliability decision outputs while preserving operator authority and the variability that *is* the operator's adaptive intelligence. Design-stage; pending empirical validation.

Eight operational integrity constraints

Architectural commitments designed to be implemented as enforceable safeguards in deployment builds — not as runtime policy. Disabling any constraint should require rebuilding from source, not toggling a flag.

1. **No coercion, ever** — the system issues recommendations, never compels.
2. **Individual baseline only** — measurements against the operator's own baseline; never against a population norm or productivity target.
3. **No surveillance architecture** — performance-support tool, not a monitoring infrastructure.
4. **Operator authority is the invariant** — the system expands effective decision options; it never narrows or preempts them.
5. **Operational transparency** — every recommendation includes a plain-language explanation.
6. **Data sovereignty** — operator-layer data belongs to the institutional program under documented data governance.
7. **Override always available** — overrides are logged for audit but never used for adverse personnel action.
8. **Independent oversight** — operator-layer deployments require institutional ethics-board review, published consent frameworks, and periodic third-party audits.

Unifying principle

> *MILO does not predict the future. It remains viable in any future.*

The principle is falsifiable: a system whose audit trail is incomplete, whose recovery is improvised, whose adaptation drifts unboundedly, or whose operator override is policy-level rather than architectural, fails the principle.

Trademark, patent, and submission status

- **Mark.** *~MILO™* — U.S. Patent and Trademark Office Serial No. 99706004; filed March 16, 2026; intent-to-use; International Class 009 (downloadable AI software). The leading tilde disambiguates from senior MILO marks held by unrelated owners in different International Classes.
- **Patent.** Patent application pending for the underlying software architecture. Implementation may require a patent license once issued; nothing in this document or its CC BY 4.0 license on the manuscript text grants any patent license.
- **Federal submission.** Submitted to the U.S. Department of Energy under the *Genesis Mission* (Executive Order 14363, November 2025); currently under review. No acceptance or grant outcome is claimed.
- **Concept DOI.** 10.5281/zenodo.20117025 — Zenodo, persistent across versions.
- **Public reference.** https://github.com/jmontano1/milo-architecture.
- **Author contact.** Jorge Enrique Flores Montano · jmontano@jmautomated.com · ORCID iD: 0009-0003-1859-8418.