CONTROLLED RECOVERABILITY FOR FUTURE AI SYSTEMS

A Technical Overview of the Plasma Harmonic Engine Architecture

Patent-pending Plasma Harmonic Engine architecture

A research initiative of Blocra Inc.

DOI: 10.5281/zenodo.21683941

AuricSphere is advancing the next generation of AI infrastructure through plasma harmonic field modeling, simulation-led research, and controlled recoverability architecture.

Public technical overview | Version 1.0 | July 2026

Contents and Research Status

This paper is organized as a research-oriented public technical overview rather than a marketing brochure. It describes the filed architecture, the current conceptual simulation, the active modeling direction, the evidence standard for future validation, and the limitations that remain.

Research maturity at publication

Filed. Patent-pending architecture describing controlled plasma harmonic recoverability of AI state capsules.

Public. Conceptual browser simulation that illustrates recovery windows, thresholds, field conditions, lifecycle state, and receipt logic.

Active. Simulation-led research toward reduced-order models, sensitivity analysis, parameter sweeps, and reproducible tests.

Planned. Higher-fidelity plasma modeling, expert plasma-physics review, chamber studies, and eventual laboratory validation.

Executive Summary

AI state is becoming strategic infrastructure. Long-running agents, orchestration systems, autonomous planners, and distributed compute workflows generate compact internal state - checkpoint digests, memory commitments, trace digests, safety records, calibration states, and continuity records - that future systems may need to resume tasks, prove continuity, audit decisions, or preserve a bounded record of a mission. Surveys of large-language-model-based autonomous agents describe the growth of agent architectures, while dedicated work on agent memory identifies memory as a key component for systems operating across long-term, complex interactions [2,3].

Conventional digital storage was built primarily for durable persistence. Deletion and expiry are usually implemented through software policy, media sanitization, key destruction, access revocation, or timers. NIST guidance treats media sanitization as a process for rendering access to target data infeasible for a given level of effort, while separate NIST guidance addresses the management and protection of cryptographic keying material [5,6]. These mechanisms are useful, but they do not make recoverability itself depend on the physical evolution of the medium that represents the state.

AuricSphere proposes a different primitive: controlled recoverability. In the patent-pending Plasma Harmonic Engine (PHE) architecture, a compact AI state capsule is mapped into harmonic behavior of a physical plasma field or a physics-based plasma field model. As the field evolves, the encoded structure may degrade in measurable ways. A recoverability score, R(t), is derived from decoder or verifier output together with measured or simulated field behavior. Recovery or verification is permitted only while two conditions hold at once: a cryptographic access condition and a physical field recoverability condition [1].

The architecture is designed around a recovery window rather than indefinite persistence. While the window remains open, a capsule may be recovered or verified if access is also authorized. As field evolution reduces recoverability below a threshold, the system may deny recovery, record expiry, and generate a recoverability receipt describing the relevant modeled or measured conditions. The filed application also describes calibration embodiments in which field parameters and excitation settings are adjusted toward a target recovery window or recovery-window class [1].

This document is a public technical overview of a filed architecture and an active research program. It explains the PHE concept, the public conceptual simulator, and the staged path toward higher-fidelity modeling and future physical validation. It does not claim that a physical plasma recoverability chamber has been built, that physical recoverability has been experimentally demonstrated, or that any specific plasma parameter set has already been validated to produce a target recoverability lifetime.

1. The Problem: AI State Is Becoming Strategic Infrastructure

Artificial intelligence systems increasingly operate as long-running agents, model orchestration systems, scientific assistants, autonomous planning systems, simulation controllers, and hybrid classical-quantum workflows. Surveys of LLM-based autonomous agents describe systems that combine planning, memory, tool use, and environment interaction across multiple application domains [2]. Dedicated work on agent memory identifies memory as a key module because it allows agents to accumulate experience and retrieve useful information across extended interactions [3].

Such systems generate internal state that may include model checkpoint digests, latent memory vectors, tool-use records, inference trace digests, safety state records, calibration states, and continuity records. These representations do not necessarily need to contain a full model or a full user-data record. In many cases, the operationally important object may be a compact commitment, digest, state fragment, or continuity record that is sufficient to resume, verify, audit, or constrain a future operation.

The importance of compact state representations grows when systems must satisfy several goals at once:

resume interrupted tasks and prove continuity across sessions, machines, or distributed nodes;

audit decision history within a bounded interval without retaining every underlying artifact indefinitely;

isolate suspicious, anomalous, or safety-relevant states for temporary review;

preserve mission-bounded or experiment-bounded continuity while limiting indefinite recoverability.

The last requirement is poorly served by persistence-first infrastructure. Continuity and auditability pull toward retention; privacy, safety, governance, and operational compartmentalization pull toward bounded lifetimes. The NIST AI Risk Management Framework emphasizes governance, mapping, measurement, and management of AI risks across the AI system lifecycle, including accountability, transparency, safety, security, resilience, and privacy-related characteristics [4].

Distributed and non-terrestrial settings

The tension becomes sharper in edge, orbital, spacecraft, and distributed-compute settings where communication delay, local power and thermal constraints, changing topology, intermittent connectivity, and node trust status complicate centralized lifecycle enforcement. Space-edge and onboard-AI literature likewise frames local processing as a response to latency, resource constraints, and trustworthiness requirements [9,10]. A mechanism whose recoverability state can be evaluated locally, rather than only by consulting a remote policy service, may therefore have distinct infrastructure value in future distributed systems.

2. Why Conventional Storage Is Not Enough

Conventional digital storage systems - cloud object stores, solid-state drives, vector databases, ledgers, and encrypted memory systems - are excellent at durable persistence. Expiry is typically implemented through key deletion, access revocation, timer-based deletion, overwriting, retention-policy enforcement, media sanitization, or cryptographic erasure. NIST SP 800-88 Rev. 2 defines media sanitization as rendering access to target data on media infeasible for a given level of effort and provides program-level guidance for sanitization controls [5]. NIST SP 800-57 provides complementary guidance for cryptographic key management [6].

These approaches are useful, and the PHE architecture does not position them as inadequate for their intended roles. They share a structural property, however: the memory lifecycle is generally separate from the ongoing physical dynamics of the storage substrate. A record is removed because a policy instructs deletion, a credential is revoked, media is sanitized, or a key is destroyed. Recoverability is not ordinarily treated as a continuously measurable consequence of the medium's physical evolution.

Software-governed memory systems can add increasingly sophisticated retention schedules, sharing rules, audit logs, revocation controls, learned forgetting, or memory-management policies. Those controls can be important, but unless recoverability itself is tied to measurable substrate behavior, the lifecycle remains fundamentally a policy layer imposed on top of a persistence-oriented medium.

The Plasma Harmonic Engine explores a different design question: can the evolving physical state of a plasma field, or a physics-based model of that field, become an independent condition on whether a compact AI state representation remains recoverable? This does not replace cryptography or ordinary storage. It adds a second dimension: authorization governs who may attempt recovery; field-derived recoverability governs whether recovery remains physically or physically modeled as possible at all.

3. The AuricSphere Thesis

Cryptographic systems can decide who may access a state. AuricSphere explores whether the ability to recover or verify a state can also be bounded by measurable physical or physics-modeled field conditions. Three ideas follow from this thesis:

Access is necessary, but not sufficient. A valid credential may be required for recovery, but authorization alone does not reopen a window once the field no longer supports the encoded structure.

Recoverability becomes a lifecycle property. Instead of a binary stored/deleted state, a capsule has a recoverability trajectory that can condition recovery, refresh, quarantine, fork, audit, sealing, and expiry.

Expiry can be evidenced rather than merely asserted. A threshold crossing can be associated with measured or modeled field conditions and recorded as a lifecycle event.

The resulting infrastructure concept is controlled recoverability: a system in which a compact AI state is intended to remain recoverable for a measured or modeled interval and to become non-recoverable after field evolution carries the encoded structure below a defined threshold. The threshold may represent reconstruction accuracy, verification confidence, harmonic coherence, decoder confidence, or another disclosed recoverability measure [1].

4. Plasma Harmonic Engine Architecture

The patent-pending PHE architecture is a multi-stage pipeline. Each stage described here is disclosed at a high level in AuricSphere's filed patent application; this public paper intentionally avoids implementation-specific recipes that are unnecessary to explain the architecture [1].

AI State Capsule

A compact representation of machine state - for example, a model checkpoint digest, latent state fragment, agent memory state, inference trace digest, tool-use record, safety state record, calibration state, continuity record, or another compact machine-state representation. The capsule need not contain a full model, full user-data record, or full reasoning transcript [1].

Access Value

A cryptographic or access value that functionally binds the capsule to encoding or verification conditions. In disclosed embodiments, the value may influence harmonic channel assignment, phase schedule, interleaving, decoder initialization, a plasma memory coordinate, feature alignment, or related conditions. A nonmatching access value can therefore prevent recovery or verification even while the field remains physically recoverable [1].

Plasma Harmonic Encoding

The capsule and access value are mapped into a plasma harmonic modulation sequence. At an architectural level, the sequence may specify field-excitation relationships involving frequencies, phases, amplitudes, channels, spatial patterns, pulse timing, or temporal windows. Certain embodiments also allow redundancy or error correction so that partial degradation does not immediately eliminate recoverability within an intended window [1].

Field Evolution

The modulation sequence is applied to a physical plasma field or a physics-based plasma field model. The field may evolve under dynamics including wave dispersion, nonlinear mode coupling, turbulence, damping, spectral entropy growth, phase decoherence, collisional effects, or loss of harmonic coherence. Standard plasma-physics references describe wave, diffusion, instability, damping, nonlinear, and turbulent behavior that motivates this class of physical model [7,8].

Harmonic Feature Extraction

Diagnostic features are derived from measured plasma sensor data or physics-simulation data representing the same dynamics. The disclosed feature space may include harmonic coefficients, phase coherence, cross-channel coherence, spectral entropy drift, damping behavior, field-energy distribution, recurrence measures, spatial modes, or other plasma-diagnostic representations [1].

Decoder or Verifier

A decoder or verifier processes the evolved harmonic feature representation. Its role is to determine whether the encoded capsule can still be reconstructed or verified under the appropriate access-conditioned configuration. The filed application permits several classes of decoder or verifier rather than fixing one public implementation [1].

Recoverability Score and Lifecycle Gate

A recoverability score is generated from the recovery or verification output together with at least one measured, sensor-derived, or physics-simulation-derived property of the evolved field. That score is compared with a recoverability threshold. Recovery, verification, or an access-dependent lifecycle operation is permitted only when both the cryptographic access condition and the physical field recoverability condition are satisfied [1].

Implementation modes

The filed architecture may be implemented as software operating on a physics-based plasma field model, as a physical plasma recoverability chamber, or as a combined simulation-and-hardware platform [1]. The public AuricSphere website currently represents the conceptual simulation path only. It should therefore be read as an explanatory model of architecture and lifecycle logic, not as a claim that the physical chamber embodiment has already been constructed or validated.

5. Controlled Recoverability and R(t)

R(t) is a recoverability score over time: a quantitative measure of the degree to which an encoded AI state capsule remains recoverable or verifiable from the evolved field. It is derived from decoder or verifier output together with measured or physics-simulation-derived properties of the evolved plasma harmonic field [1]. The filed application permits scalar, vector, score-curve, threshold-classification, probability, or confidence-based forms rather than requiring one fixed public formula.

R(t) = f(verification confidence, harmonic coherence, phase stability, spectral behavior, decoder confidence, ...)

The expression above is intentionally abstract. It communicates the role of the score without publishing a single production weighting rule. The patent application expressly permits multiple scoring forms and feature combinations [1]. A public whitepaper therefore does not need to prescribe exact score weights, calibration coefficients, feature normalization rules, or decoder-specific thresholds.

The recovery window is the interval during which R(t) remains at or above a selected recoverability threshold. T-half is a recoverability half-life descriptor; Tmax is the earliest time at which R(t) falls below the threshold and remains there under a defined tolerance condition [1]. These descriptors turn recoverability into a time-dependent system property rather than a static binary flag.

Different plasma conditions and encoding settings are expected to produce different modeled or measured recoverability trajectories. In the filed calibration embodiments, a recovery-window controller adjusts field parameters and excitation settings toward a target recovery window, target class, half-life range, coherence budget, or score-curve envelope [1]. This is an architectural capability disclosed in the filing; it is not represented here as an experimentally demonstrated result.

In a physical embodiment, expiry is intended to arise because plasma-field evolution has degraded the encoded harmonic structure to the point that the field no longer satisfies the recoverability condition. That is conceptually distinct from deletion by instruction alone. The scientific question is whether such degradation can be made sufficiently measurable, repeatable, and controllable to support useful recovery-window classes.

6. Dual-Condition Recovery

Recovery, verification, and access-dependent lifecycle operations are permitted only when two independent conditions are satisfied simultaneously: a cryptographic access condition and a physical field recoverability condition. The latter is based on measured or physics-simulation-derived field behavior and is not satisfied by software policy, key availability, or timer state alone [1].

The decision logic is straightforward. If the access condition is valid and the field recoverability condition remains above threshold, recovery or verification may be permitted. If access is valid but the field condition has failed, recovery is denied and expiry may be recorded. If access is invalid, recovery is denied regardless of whether the field itself remains recoverable. Once both access and field recoverability are absent, the capsule is both unauthorized and non-recoverable within the modeled lifecycle state.

This differs from ordinary authorization in a specific way: a valid credential cannot, by itself, recover a capsule after the field-derived recoverability condition has failed. The filed application describes this as physically permissioned memory continuity [1]. The physical or physics-modeled state of the medium becomes a second gate rather than merely a passive location from which authorized software retrieves data.

This separation also clarifies the role of cryptography. The PHE architecture does not eliminate cryptographic access control. Instead, cryptographic access and physical recoverability perform different functions. One answers whether an actor or system is authorized to attempt recovery. The other answers whether the evolved representation still contains sufficient recoverable structure for that attempt to succeed under the disclosed conditions.

7. Lifecycle Operations and Recoverability Receipts

The dual-condition gate governs a family of lifecycle operations, not retrieval alone. Disclosed operations include recovery, verification, refresh, quarantine, fork, audit, promotion, demotion, sealing, expiry, and receipt generation [1]. This makes the architecture a state-lifecycle system rather than a one-time storage primitive.

Recover and verify. reconstruct or confirm a capsule while the required access and field conditions remain satisfied.

Refresh. re-encode a still-recoverable capsule, potentially restarting or changing its recovery window under an allowed lifecycle policy.

Quarantine. place a suspicious or anomalous capsule into a more restrictive recoverability class or shorter review window.

Fork. encode a common capsule into multiple states or coordinates associated with different target windows, supporting rollback, experimentation, review, or controlled degradation research.

Audit. permit bounded inspection while the required access and field conditions remain satisfied.

Seal and expire. deny further recovery after a threshold crossing or lifecycle denial event and record the resulting state.

Recoverability receipts

A recoverability receipt is an evidence artifact - a record, certificate, attestation-style object, log entry, or evidence package - indicating what happened to a capsule's recovery status and under what field conditions. Public-level fields may include a capsule identifier, recovery-window identifier, score or score-curve digest, threshold-crossing event, field-condition summary, decoder or verifier confidence, sensor or simulation provenance, lifecycle operation, expiry status, timestamp, and content hash [1].

Recoverability receipts may be generated in association with recovery, refresh, quarantine, fork, audit, threshold crossing, or physical expiry. Their purpose is auditability: a reviewing party can inspect the field or model evidence associated with the lifecycle event rather than relying solely on a declaration that deletion or expiry occurred. Receipts in this architecture are records; this paper makes no blockchain or legal-evidence claim about them.

In a future physical implementation, the evidentiary value of a receipt would depend on the quality of the underlying measurements, sensor provenance, calibration provenance, timestamping, signing or hashing process, and the validity of the recoverability model. A receipt should therefore be treated as an output of a measurement and lifecycle system, not as a substitute for validating that system.

8. Conceptual Simulation

AuricSphere's public website includes an interactive conceptual simulator of the PHE architecture. It visualizes a capsule's modeled recoverability over time: an R(t) curve against a selectable threshold, field conditions responding to user controls, independent access and physical-field conditions, lifecycle transitions from recoverable through degrading to expired, and receipt generation at threshold crossing.

The purpose of the public simulator is legibility. Controlled recoverability is easier to evaluate when a viewer can observe how a recovery window, threshold, lifecycle gate, and expiry event relate to one another. The simulator also provides a common visual vocabulary for future research discussions without requiring publication of implementation-specific calibration methods.

The simulator is intentionally limited. It is not a laboratory validation, a production plasma chamber, a digital twin of a validated chamber, or proof that physical plasma parameters already produce the displayed recovery trajectories. Its controls should be understood as conceptual model variables that make the disclosed architecture inspectable. Any future claim that particular physical plasma parameters program a particular recoverability lifetime would require separate modeling and measurement evidence.

A useful public simulator should therefore satisfy two standards at once: it should be faithful to the filed architecture, and it should be explicit about the difference between architecture visualization and scientific validation. AuricSphere treats that distinction as part of the research program rather than as a weakness to hide.

9. Research Roadmap

AuricSphere is pursuing a staged, simulation-led path from filed architecture toward falsifiable physical tests. Later stages are planned, not complete. The objective is not to decorate the concept with increasingly complex models; it is to identify where the architecture succeeds, fails, or requires revision before higher-cost physical work.

Stage 1 - Filed architecture and technical framing

Define the controlled recoverability primitive, capsule lifecycle, physical field recoverability condition, scoring concepts, calibration embodiments, and physical chamber embodiment in a patent-pending architecture.

Stage 2 - Public conceptual simulation

Make the core architecture legible through a browser model that shows R(t), threshold behavior, dual-condition gating, lifecycle state, and receipt logic without representing laboratory validation.

Stage 3 - Reduced-order modeling

Express the disclosed architecture as inspectable equations, physically motivated state variables, reproducible parameter sets, and deterministic or seeded simulation experiments.

Stage 4 - Sensitivity analysis and parameter sweeps

Identify which modeled plasma variables most strongly influence recoverability trajectories and whether their effects remain directional, separable, and stable across seeds and parameter regions.

Stage 5 - Higher-fidelity plasma modeling

Translate the most promising reduced-order relationships into more physically complete numerical models, including appropriate plasma-wave, damping, coupling, turbulence, and diagnostic representations.

Stage 6 - Expert review and chamber studies

Subject assumptions, observables, parameter ranges, scaling choices, and proposed diagnostics to independent plasma-physics scrutiny before committing to a physical test program.

Stage 7 - Future laboratory validation

Test whether measured plasma-field behavior can produce recoverability trajectories that are repeatable, calibratable, and predictive enough to support physically governed recovery-window classes.

Progress through this roadmap is conditional. A later stage should be entered only when the preceding stage provides evidence strong enough to justify the additional complexity or cost. Negative results are informative: they can identify parameter regions, scoring assumptions, or embodiments that should be narrowed or abandoned.

10. Near-Term Validation Program

The central near-term research question is whether field parameters can govern recoverability lifetime in a repeatable, model-grounded way. The filed architecture already describes calibration loops that tune or characterize parameters toward target recovery windows [1]. The research task is now to determine whether that relationship survives increasingly realistic modeling and, eventually, measurement.

10.1 Sensitivity

Changing a physically meaningful parameter should produce a directional and quantifiable change in the recoverability trajectory, not merely an arbitrary software response. A useful sensitivity study should report effect size, direction, uncertainty, and the parameter range over which the effect is observed.

10.2 Repeatability

The same parameter region should produce comparable recovery-window behavior across controlled trials or simulation seeds within stated uncertainty. If small changes in seed or numerical settings produce large uncontrolled changes in Tmax or T-half, the parameter relationship is not yet suitable for calibration.

10.3 Separability

The effect of one parameter should be distinguishable, where possible, from confounding changes in noise, decoder configuration, redundancy, or unrelated field variables. This does not require perfect isolation in a nonlinear system, but it does require enough structure to identify which variables materially influence the recoverability trajectory.

10.4 Prediction

A calibration model should predict an approximate window, half-life range, or window class for held-out conditions. Predictive performance should be reported rather than assumed. A result that fits only the training sweep but fails on held-out parameter combinations would not establish programmable recoverability.

10.5 Physical traceability

Any simulation term used to govern R(t) should map to a physical field quantity, plasma-dynamics term, diagnostic feature, or experimentally testable proxy. The research program should be able to explain what a modeled variable would correspond to in a chamber or diagnostic system.

10.6 Measurement design

Before physical testing, the proposed observables should be tied to realistic diagnostics, sampling rates, noise floors, calibration procedures, and uncertainty estimates. The physical experiment should be designed around measurable quantities rather than around variables that exist only inside a convenient software model.

10.7 Falsifiability

A credible program must allow negative outcomes. If plausible plasma regimes do not produce stable, separable, and measurable relationships between field evolution and capsule recoverability, that result should constrain or falsify specific embodiments. If such relationships do appear, the next question is whether they can be calibrated into useful recovery-window classes without relying on a software timer disguised as physics.

Evidence that would materially advance the architecture

The strongest near-term evidence would not be a visually convincing simulation alone. It would be a reproducible result showing that one or more physically meaningful plasma parameters systematically shift a recoverability trajectory; that the effect remains measurable across repeated trials; that a calibration model can predict the direction or approximate magnitude of the shift in held-out conditions; and that the same variables can be connected to realistic plasma diagnostics or higher-fidelity physics models.

Such evidence would support the proposition that recoverability lifetime can be treated as a field-governed system property rather than as an arbitrary software countdown. It would still not, by itself, constitute physical chamber validation. It would establish a more rigorous basis for deciding which physical experiment to build and what result that experiment must reproduce.

11. Strategic Relevance

Controlled recoverability is potentially relevant wherever AI state must be both recoverable and bounded. The significance is not that every AI system requires plasma-based state continuity; it is that some future systems may benefit from an independent, measurable condition on whether a compact state can still be recovered.

Frontier AI infrastructure

Long-running agents may require continuity, rollback, and bounded audit without indefinite retention of every internal state. Agent-memory literature identifies memory as central to agents' ability to accumulate experience and support future action [3]. A physically governed recovery window would add a distinct lifecycle primitive to that software memory stack.

Safety-critical and sensitive workflows

Safety-state records, quarantine paths, temporary audit states, and mission-bounded continuity may benefit from lifecycle evidence tied to measurable or modeled field conditions. NIST AI RMF 1.0 provides a widely used framing for governing and measuring AI risk across system lifecycles [4]. Controlled recoverability would not replace those governance practices; it would provide a physical or physics-modeled condition that such policies could reference.

Auditability and lifecycle governance

Recoverability receipts associate lifecycle events with field-derived or model-derived evidence. This may be relevant to systems that need to demonstrate not only that a state was declared expired, but that the recoverability condition itself crossed a defined boundary. The usefulness of such receipts depends on the validity of the underlying measurement and scoring system, which is why validation remains central to the program.

Space-based and distributed compute

The filed application describes optional embodiments in which recoverability policy or field conditions account for mission duration, communication delay, radiation exposure, power and thermal conditions, node trust status, and continuity requirements across terrestrial, orbital, spacecraft, lunar or planetary, and edge-compute contexts [1]. Space-edge and satellite-edge AI literature likewise emphasizes onboard processing, latency, trustworthiness, and resource constraints [9,10].

Strategic evaluation

For strategic evaluators, the significance lies in a filed architectural position around a proposed infrastructure primitive - physically governed recoverability - at the intersection of AI infrastructure, security architecture, and plasma-field modeling, together with a simulation-led research program designed to test that primitive rather than simply assert it.

12. Limitations and Current Status

Accuracy about maturity is part of the value of this document. The following limitations are intentional and should be read as the current boundary of what AuricSphere is claiming publicly.

Patent pending. AuricSphere's Plasma Harmonic Engine architecture is patent-pending. No representation is made that patents have been granted.

Public simulation is conceptual. It illustrates the architecture; it does not demonstrate physical plasma recovery.

Reduced-order modeling is research tooling. It supports sensitivity analysis, parameter exploration, falsification, and roadmap design. It is not laboratory validation and is not represented as such.

The chamber is a described embodiment. The physical plasma recoverability chamber is disclosed in the filed application but is not represented as a built or deployed system.

Parameter-to-lifetime programmability remains a research question. The patent discloses calibration and tuning toward target recovery windows. Public materials do not claim that a particular physical parameter map has already been experimentally validated.

No customer deployment claim. This whitepaper describes a filed architecture and research program, not a production deployment or current customer use.

No performance guarantee. The architecture has not yet established production-level reliability, latency, scale, cost, energy consumption, or physical lifetime ranges.

No claim that plasma is required for all AI memory. Controlled plasma harmonic recoverability is proposed for use cases where physically governed, bounded recoverability could add value; it is not positioned as a universal replacement for conventional storage.

Public disclosure boundary

This paper explains the filed architecture at a level intended to support technical evaluation. It does not publish future experimentally optimized parameter sets, implementation-specific calibration recipes, detailed encoder or decoder recipes, or other later-developed material that may warrant separate intellectual-property review before disclosure. That boundary allows public scientific communication without treating every future implementation detail as part of a marketing document.

13. Conclusion

Controlled recoverability reframes AI state as something whose recovery can be governed by more than permission. In the Plasma Harmonic Engine architecture, authorization determines who may attempt recovery; measurable field evolution determines whether recovery remains possible; and recoverability receipts may evidence which outcome occurred.

AuricSphere's patent-pending architecture explores a future in which AI state continuity, recovery, expiry, and auditability are bounded by both authorization and measurable field-derived recoverability - a primitive conventional storage was not designed to provide, for a class of future AI systems that may increasingly need bounded rather than indefinite continuity.

The next research threshold is concrete: determine whether physically meaningful plasma parameters can predictably shape recoverability lifetime in simulation and, ultimately, in measurement. That question is testable. The answer will determine how far the architecture can move from filed concept toward validated infrastructure.

References

[1] Blocra Inc. / AuricSphere. Systems and Methods for Controlled Plasma Harmonic Recoverability of Artificial Intelligence State Capsules. Patent application, patent pending.

[2] Wang, L., Ma, C., Feng, X., Zhang, Z., Yang, H., Zhang, J., Chen, Z., Tang, J., Chen, X., Lin, Y., Zhao, W. X., Wei, Z., & Wen, J.-R. (2024). A survey on large language model based autonomous agents. Frontiers of Computer Science, 18, 186345. https://doi.org/10.1007/s11704-024-40231-1.

[3] Zhang, Z., Dai, Q., Bo, X., Ma, C., Li, R., Chen, X., Zhu, J., Dong, Z., & Wen, J.-R. (2025). A Survey on the Memory Mechanism of Large Language Model-based Agents. ACM Transactions on Information Systems, 43(6), Article 155, 1-47. https://doi.org/10.1145/3748302.

[4] National Institute of Standards and Technology. (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0). NIST AI 100-1. https://doi.org/10.6028/NIST.AI.100-1.

[5] Chandramouli, R., & Hibbard, E. (2025). Guidelines for Media Sanitization. NIST Special Publication 800-88 Revision 2. https://doi.org/10.6028/NIST.SP.800-88r2.

[6] Barker, E. (2020). Recommendation for Key Management: Part 1 - General. NIST Special Publication 800-57 Part 1 Revision 5. https://doi.org/10.6028/NIST.SP.800-57pt1r5.

[7] Chen, F. F. (2016). Introduction to Plasma Physics and Controlled Fusion (3rd ed.). Springer. https://doi.org/10.1007/978-3-319-22309-4.

[8] Chian, A. C.-L. (2022). Nonlinear dynamics in space plasma turbulence: temporal stochastic chaos. Reviews of Modern Plasma Physics, 6, Article 23. https://doi.org/10.1007/s41614-022-00095-z.

[9] Gardill, M., Kinsner, W., Budroweit, J., Al-Hourani, A., Dunkel, E., Swope, J., Evans, D., & Staebler, M. (2023). Towards Space Edge Computing and Onboard AI for Real-Time Teleoperations. 2023 IEEE 22nd International Conference on Cognitive Informatics & Cognitive Computing (ICCI*CC), Palo Alto, California. IEEE CS Press.

[10] Shi, Y., Zhu, J., Jiang, C., Kuang, L., & Letaief, K. B. (2025). Satellite Edge Artificial Intelligence with Large Models: Architectures and Technologies. arXiv:2504.01676. https://doi.org/10.48550/arXiv.2504.01676.