

Deterministic Replay Architectures for AI Systems

Runtime lineage, reconstructable transitions, observability, drift detection, and replay-aware execution governance.

Publication series	Bluehand Research Library
Artifact type	Research Brief
Authority class	canonical_public
Publication status	canonical
Artifact ID	BH-RL-2026-0006
Version	1.0.0
Publication date	2026-05-19
Canonical HTML	https://www.blue-hand.org/research/deterministic-replay/
Canonical PDF	https://www.blue-hand.org/bluehand-deterministic-replay-architectures-2026.pdf
Primary route	/research/deterministic-replay/
Domains	deterministic replay, AI observability, runtime governance, execution traceability
Keywords	deterministic replay, AI observability, runtime lineage, replayable AI systems, execution traceability, drift detection

FRONT MATTER

Abstract

A Bluehand research artifact on deterministic replay as a trust and observability pattern for AI systems: runtime transitions should be reconstructable, bounded, and reviewable.

Provenance

Source: research/registry.json. HTML route /research/deterministic-replay/ remains the canonical indexed surface.

Compile receipt

Generated at 2026-06-13T02:11:50Z by scripts/pdf/generate_encodings.py (engine=reportlab; depth=published_reportlab_projection_v2).

Publication note

Portable PDF encoding of /research/deterministic-replay/. HTML is canonical; this file supports sharing, review, archive, and stakeholder packets.

DOCUMENT BODY

Purpose

Orient readers to replay as a trust mechanism for AI systems: consequential transitions should be reconstructable enough to inspect.

Problem addressed

Agentic workflows can create opaque sequences of model calls, tool use, state changes, summaries, and actions that are hard to explain after the fact.

Why now

AI systems are moving from isolated chat interactions toward persistent agents, retrieval systems, personal workflows, and institution-facing automation.

Inspiration

CP0/BIS-CP0 transition discipline, SigMem0 replay concepts, and Bluehand's broader concern with reconstructable reasoning and state mutation.

Vision direction

AI runtimes where important transitions can be replayed, audited, compared, or challenged without pretending hidden model reasoning is fully accessible.

Reader takeaway

Bluehand sees replayability as part of governance, trust, and systems maturity.

Operational thesis

A system that cannot reconstruct consequential transitions cannot reliably govern them. Deterministic replay architectures aim to preserve enough event structure, state transition context, and lineage to inspect what occurred and why.

Why it matters

Agentic workflows introduce operational ambiguity: a task can be delegated, routed, transformed, summarized, or executed through multiple intermediaries. Replayability helps convert opaque activity into inspectable execution history.

Governance boundary

This artifact defines replay as an architecture and evaluation principle. It should be integrated with implementation evidence only when actual replay logs, tests, or runtime traces are linked.

Recruiter relevance

Provides a capability signal for AI infrastructure, governance, semantic systems, agentic workflow, or local-first execution roles.

Grant relevance

Supports public-interest framing where responsible AI, trustworthy infrastructure, human-centered systems, or research-to-venture pathways matter.

Partner relevance

Helps potential partners understand the problem Bluehand is orienting around and where collaboration could fit.

Technical reviewer relevance

Surfaces the relevant problem, methods, constraints, failure modes, and implementation boundaries without requiring internal Bluehand context.

Maturity note

Published PDF; HTML should mark replay as a design/research artifact unless implementation evidence is linked.

Uncertainty boundaries

Summaries, embeddings, retrieval projections, and rendered pages are not evidentiary authority without declared lineage to source artifacts. Implementation claims require separate evidence.

Public claims

- Bluehand identifies deterministic replay as a key design pattern for trustworthy AI execution.

Uncertain claims

- Replay completeness depends on implementation, privacy constraints, and the kind of system being observed.

Implementation boundary

This is a public Research Object. Implementation evidence, strict lineage, and runtime proof belong in project/repo-specific surfaces unless explicitly linked.

Publication doctrine

The PDF is the shareable artifact, but the HTML route is the authority surface. It gives crawlers context, gives recruiters a readable summary, and gives Bluehand a stable research landing page that can rank for the work rather than the file format.

Source registry

Canonical registry object: research/registry.json (BH-RL-2026-0006). Public discovery:
<https://www.blue-hand.org/research/registry.public.json>

Lineage

```
id=BH-RL-2026-0006; source=research/registry.json; route=/research/deterministic-replay/; engine=reportlab;  
generated_by=scripts/pdf/generate_encodings.py; pdf_depth_status=published_reportlab_projection_v2
```

Do not infer

- Do not infer total access to hidden model reasoning or perfect reconstruction of every internal state.