

Personal Corpus Model Training

Lineage-Derived Personal Corpus Model Training (Closed Loop) is a secondary inventive step of the Agent-Resident Execution Substrate disclosed in U.S. Provisional Application No. 64/070,239: a model fine-tuned on user-authored and user-curated artifacts under governance policy, with lineage tracking from training data through derived outputs.

Personal Corpus Model Training

Model fine-tuned on user-authored and user-curated artifacts under governance policy, with lineage tracking from training data through derived outputs. The agent's knowledge is grounded in the user's own work rather than general public datasets.

The training corpus is assembled from the lineage field by selecting artifacts admissible under the corpus policy object, filtering for modality compatibility, and applying any declared redaction rules. A parameter-efficient fine-tuning operation updates a bounded subset of the model's parameters so the operation is feasible within the device's local compute envelope. The user's accumulated body of work is internalized in the parameter values rather than consulted as external context at inference time.

Disclosure Scope

This article describes subject matter disclosed in U.S. Provisional Application No. 64/070,239. As a provisional disclosure, the specific mechanisms are described in illustrative embodiments and may differ in any nonprovisional application claiming priority to it.

Agent-Resident Execution

[All 40 steps → \(/inventive-steps\)](#)

Substrate (/agent-resident-execution-substrate)

Persistent execution environment carried by the agent, not the host — identity, state, and lineage across power cycles, devices, and upgrades.

Provisional application

PRIMARY TECHNICAL DISCLOSURE

- [Agent-Resident Execution Substrate, Articles \(/articles/agent-resident-execution-substrate\)](#)

SECONDARY TECHNICAL

- [Persistent Semantic Agent \(/articles/agent-resident-execution-substrate/persistent-semantic-agent\)](#)
- [Managed Inference Tool Registry \(/articles/agent-resident-execution-substrate/managed-inference-tool-registry\)](#)
- [Agent-to-Tool Dispatcher \(/articles/agent-resident-execution-substrate/agent-to-tool-dispatcher\)](#)
- [Lineage-Derived Training Signal \(/articles/agent-resident-execution-substrate/lineage-derived-training-signal\)](#)
- [Identity Preservation Across Upgrades \(/articles/agent-resident-execution-substrate/identity-preservation-across-upgrades\)](#)
- [Cognitive State-Conditioned Dispatch \(/articles/agent-resident-execution-substrate/cognitive-state-conditioned-dispatch\)](#)
- [Governed Tool Lifecycle \(/articles/agent-resident-execution-substrate/governed-tool-lifecycle\)](#)
- [Continuity-Proof Lineage \(/articles/agent-resident-execution-substrate/continuity-proof-lineage\)](#)
- [Substrate Runtime Continuity \(/articles/agent-resident-execution-substrate/substrate-runtime-continuity\)](#)

- [Personal Corpus Model Training \(/articles/agent-resident-execution-substrate/personal-corpus-model-training\)](/articles/agent-resident-execution-substrate/personal-corpus-model-training).
- [Heterogeneous Inference Endpoints \(/articles/agent-resident-execution-substrate/heterogeneous-inference-endpoints\)](/articles/agent-resident-execution-substrate/heterogeneous-inference-endpoints).
- [Atomic Lifecycle Substitution \(/articles/agent-resident-execution-substrate/atomic-lifecycle-substitution\)](/articles/agent-resident-execution-substrate/atomic-lifecycle-substitution).
- [Integrity Signal Feedback \(/articles/agent-resident-execution-substrate/integrity-signal-feedback\)](/articles/agent-resident-execution-substrate/integrity-signal-feedback).
- [Hardware-Bound Identity \(/articles/agent-resident-execution-substrate/hardware-bound-identity\)](/articles/agent-resident-execution-substrate/hardware-bound-identity).
- [Cognitive State Append-Only Invariant \(/articles/agent-resident-execution-substrate/cognitive-state-append-only-invariant\)](/articles/agent-resident-execution-substrate/cognitive-state-append-only-invariant).
- [Counterparty Identity Records \(/articles/agent-resident-execution-substrate/counterparty-identity-records\)](/articles/agent-resident-execution-substrate/counterparty-identity-records).
- [Privacy Egress-Controlled Disclosure \(/articles/agent-resident-execution-substrate/privacy-egress-controlled-disclosure\)](/articles/agent-resident-execution-substrate/privacy-egress-controlled-disclosure).
- [Federated Cross-Device Agent Identity \(/articles/agent-resident-execution-substrate/federated-cross-device-agent-identity\)](/articles/agent-resident-execution-substrate/federated-cross-device-agent-identity).

APPLICATIONS · GENERAL

- [Personal AI Agents That Survive Device Loss: One Continuous Identity and a Private Corpus Across Every Device \(/articles/agent-resident-execution-substrate/personal-cross-device-agents\)](/articles/agent-resident-execution-substrate/personal-cross-device-agents).
- [Enterprise Agent Fleets: Stable Agent Identity and Governed Tool Access Across Model Upgrades and Infrastructure Migration \(/articles/agent-resident-execution-substrate/enterprise-agent-fleets\)](/articles/agent-resident-execution-substrate/enterprise-agent-fleets).
- [Audit-Grade Agent Identity for Regulated Finance and Healthcare: Continuity-Proof Lineage Across the Agent Lifecycle \(/articles/agent-resident-execution-substrate/regulated-industry-agents\)](/articles/agent-resident-execution-substrate/regulated-industry-agents).
- [Edge and On-Device Agents: Hardware-Bound Identity Across Heterogeneous Inference Endpoints \(/articles/agent-resident-execution-substrate/edge-and-on-device-agents\)](/articles/agent-resident-execution-substrate/edge-and-on-device-agents).
- [Agent-to-Agent Commerce With Counterparty Identity Records and Egress-Controlled Disclosure \(/articles/agent-resident-execution-substrate/agent-to-agent-commerce\)](/articles/agent-resident-execution-substrate/agent-to-agent-commerce).
- [Governed Tool Lifecycles for Managed Inference-Provider Ecosystems: A Substrate Approach to Owning, Routing, and Retiring AI Tools \(/articles/agent-resident-execution-substrate/managed-to-ol-ecosystems\)](/articles/agent-resident-execution-substrate/managed-to-ol-ecosystems).
- [Proving Unbroken Continuity in Long-Lived Autonomous Systems Across Substrate Migration and Atomic Model Substitution \(/articles/agent-resident-execution-substrate/long-lived-autonomous-systems\)](/articles/agent-resident-execution-substrate/long-lived-autonomous-systems).

- [Personal-Model Personalization: A User's Own Corpus-Internalized Model on the Agent-Resident Execution Substrate](/articles/agent-resident-execution-substrate/personal-model-personalization) (/articles/agent-resident-execution-substrate/personal-model-personalization).
- [On-Device Agent Identity for Robots and Autonomous Vehicles: An Auditable Substrate for Embodied Physical-World Agents](/articles/agent-resident-execution-substrate/embodied-physical-world-agents) (/articles/agent-resident-execution-substrate/embodied-physical-world-agents).

APPLICATIONS · SPECIFIC

- [LangGraph Platform \(LangChain\) vs an agent-resident execution substrate: orchestration-graph state versus a portable, hardware-anchored agent runtime](/articles/agent-resident-execution-substrate/langgraph-platform) (/articles/agent-resident-execution-substrate/langgraph-platform).
- [OpenAI AgentKit and the Assistants/Responses API vs agent-carried, hardware-anchored identity with governed tool lifecycle](/articles/agent-resident-execution-substrate/openai-agentkit) (/articles/agent-resident-execution-substrate/openai-agentkit).
- [Microsoft Copilot Studio vs an agent-resident execution substrate: platform-hosted agent authoring versus portable, device-resident agent identity and continuity](/articles/agent-resident-execution-substrate/microsoft-copilot-studio) (/articles/agent-resident-execution-substrate/microsoft-copilot-studio).
- [Google Vertex AI Agent Engine \(managed runtime for deploying and scaling agents, with sessions/memory\) vs an agent-carried, continuity-proofed identity substrate](/articles/agent-resident-execution-substrate/google-vertex-agent-engine) (/articles/agent-resident-execution-substrate/google-vertex-agent-engine).
- [AWS Bedrock AgentCore \(runtime, memory, identity, and gateway services for deploying agents at scale\) vs an agent-resident execution substrate: where does the agent identity actually live?](/articles/agent-resident-execution-substrate/aws-bedrock-agentcore) (/articles/agent-resident-execution-substrate/aws-bedrock-agentcore).
- [Letta \(formerly MemGPT\) vs an append-only cognitive-state substrate: what a memory-management framework does not provide](/articles/agent-resident-execution-substrate/letta-memgpt) (/articles/agent-resident-execution-substrate/letta-memgpt).
- [Cognition's Devin, an autonomous AI software-engineering agent vs a portable, continuity-proofed agent-resident runtime](/articles/agent-resident-execution-substrate/cognition-devin) (/articles/agent-resident-execution-substrate/cognition-devin).
- [Cloudflare Agents \(Durable Objects\) vs an agent-resident execution substrate: portable hardware-bound identity and continuity-proof lineage](/articles/agent-resident-execution-substrate/cloudflare-agents) (/articles/agent-resident-execution-substrate/cloudflare-agents).
- [Ollama alternative: from local model runner to a governed agent-resident substrate](/articles/agent-resident-execution-substrate/ollama) (/articles/agent-resident-execution-substrate/ollama).
- [Apple Intelligence \(on-device foundation models, Private Cloud Compute\) vs a persistent agent-resident substrate: who owns identity, lineage, and the model?](/articles/agent-resident-execution-substrate/apple-intelligence) (/articles/agent-resident-execution-substrate/apple-intelligence).

[Agent-Resident Execution Substrate overview](/agent-resident-execution-substrate) → (/agent-resident-execution-substrate)

