Nick Clark · Published April 25, 2026 · PDF

Personal AI Agents That Survive Device Loss: One Continuous Identity and a Private Corpus Across Every Device

When a phone is lost, stolen, or upgraded, today's personal AI assistant resets to a blank slate, because the assistant is a stateless cloud session keyed to an account rather than a persistent entity that the user owns. The Agent-Resident Execution Substrate, disclosed in U.S. Provisional Application No. 64/070,239, resolves this by making the agent itself the durable object: a hardware-bound identity, an append-only history under a continuity proof, and a personal corpus model whose weights internalize the user's own work, all migratable to a new device under a governed operation. It draws on sibling portfolio inventions for memory-native identity and authentication and for cognition-compatible network transport to carry that identity across a user's device population.

What This Application Specifies

This application describes how a single personal AI agent can follow a person across all of their devices as one continuous identity, carrying a private corpus of that person's accumulated work, and survive losing or replacing any individual device. The agent is not a chat session and not a cloud account profile. It is a persistent execution substrate

that runs on the device and owns the things that make it personal: a persistent identity field, a cognitive state field, an append-only lineage field of everything it has done, and a governance policy field that the user controls.

The substrate, as disclosed in U.S. Provisional Application No. 64/070,239, treats inference models as managed assets subordinate to the agent rather than as the agent itself. A managed inference tool registry holds one or more local model endpoints (a general writing model, a code model, a personal corpus model, and others), and an agent-to-tool dispatcher routes each request to the right endpoint conditioned on the agent's current cognitive state, the input modality, and the applicable policy. Because the agent's identity is separate from any particular model, the spec is explicit that the agent's identity is not dependent on any specific model artifact and is preserved across replacement, retraining, or removal of any subordinate model. The same separation is what makes a device upgrade survivable: nothing about who the agent is lives in the silicon that gets thrown away.

Two disclosed mechanisms do the cross-device work directly. First, hardware-anchored identity binding cryptographically ties the persistent identity field to a hardware security element (a secure enclave, trusted platform module, or embedded secure element) whose private key material is not extractable. Second, a governed migration operation moves that identity to a new device through attestation by the originating hardware element, key derivation by the destination hardware element under a migration policy object, and recording of the migration event in the lineage field with attestations from both devices. A device upgrade is, in this architecture, a normal governed lifecycle event, not a fresh start.

Why It Matters

A personal assistant becomes valuable in proportion to how much of your life it has absorbed: your writing voice, your recurring projects, the people you deal with, the conventions you never bother to state. That accumulated value is exactly what the

prevailing architecture cannot retain or protect. Network inference services, as the background of the application notes, maintain no persistent representation of any individual user's body of work as a property of the model itself; the model is shared across all users, and your context is supplied per request and discarded between requests. Retrieval systems bolt a document index onto that shared model, but the corpus is consulted at each query and never internalized.

So when a device is lost or replaced, there is nothing of the user's agent to carry forward, because the personalization never lived on the device as a transferable object in the first place. It lived as transient prompt context and as files in a cloud account. The user re-onboards from zero. Worse, the private material that did accumulate, the prompts and documents shared to a remote service, left the user's control the moment it was sent.

This application matters because it relocates both the value and the control to an object the user actually possesses. The personal corpus model internalizes the user's accumulated work in its weight parameters rather than consulting it as external context, so personalization is intrinsic and portable. The privacy invariant holds that lineage records, model artifacts, personal corpus model parameters, and counterparty identity records are not transmitted off the substrate device except under an explicit disclosure policy object. Losing a device stops being a loss of identity, and using the agent stops being a continuous leak of private material.

How It Composes With the Domain

Consider one user with a phone, a laptop, and a tablet. Each runs its own substrate instance with its own agent, tool registry, and lineage store, and the application provides for exactly this multi-tier deployment: the substrate is operable across at least two of a mobile, tablet, laptop, desktop, head-mounted, wearable, vehicular, or household-appliance tier, with the user's persistent identity field maintaining continuity across hardware tiers through a cross-tier policy object.

The devices are bound into one logical agent through the disclosed federation layer. Federation here is deliberately lightweight: it comprises exchange of lineage records and does not require exchange of model artifacts, so each device keeps responsibility for its own models while incorporating outcome signals observed on the others. Above that exchange sits the mechanism that makes this feel like a single assistant rather than three: a federated agent identity record that verifies through cross-device attestations that the federated agents correspond to a single user identity, so that inference requests, lineage records, scope mutations, and counterparty encounters across the devices are treated as originating from one agent. That record is preserved across device additions, device retirements, and device hardware refresh. Notably, federation operates without reliance on a centralized federation authority, coordinating through anchor-governed resolution and memory-native transport drawn from sibling portfolio applications, so the user's cross-device identity does not depend on any one vendor's cloud staying online.

Within that single identity, scope partitioning keeps contexts cleanly separated. The agent maintains named scope records (a personal scope, a work scope, a household scope, a project scope), each with an independent corpus policy, tool subset, and lineage partition, all under one persistent identity. A scope-selection rule can switch the active scope automatically on time of day, calendar event, device location, or originating application, so the laptop in the morning and the phone at a client site present the appropriate context without a separate login.

Three failure and transition modes map onto disclosed operations:

- Planned upgrade. Transfer-provisioning initializes the new device by transfer of the persistent identity field, lineage field, cognitive state field, governance policy objects, and personal corpus model artifacts from the prior device under a governed agent migration operation, with attestations recorded on both devices. The new phone wakes up as the same agent, mid-history.

- Lost or stolen device. Because the identity is hardware-bound and not transferable to a device lacking the bound hardware element except under governed migration, a thief holds an inert shell. The substrate can additionally quarantine agent execution, suspending dispatch and refusing lifecycle operations, if the hardware element fails attestation. The user restores onto replacement hardware from an agent snapshot artifact, which carries a continuity attestation linking it to the agent's identity continuity proof, through the disclosed agent restore and migration lifecycle operations.
- Guest or shared use. A guest-provisioning mode instantiates a temporary, decommissioned-on-exit agent isolated from the primary user's state under the privacy invariant, so handing someone your tablet does not expose your corpus.

When a local model genuinely cannot handle a request, cloud-burst forwarding can send it to a remote endpoint, but only after a disclosure test confirms the input is admissible for off-device disclosure, and the forwarded payload is itself recorded as a disclosure event. Reaching for more capability never silently becomes a privacy breach.

What This Enables

A personal agent that genuinely belongs to its user. Because the personal corpus model is fine-tuned against artifacts the user authors and feeds the next round of fine-tuning from the lineage of accepted and revised outputs, the agent's grasp of the user's voice and conventions improves continuously and locally, without manual corpus curation. The training signal is drawn from real downstream outcomes (acceptance, revision, execution success or failure) rather than from the model's own prior output, which the application notes mitigates the distributional collapse seen when a model trains on its own generations.

Concretely, the architecture enables: a writing or coding assistant whose understanding of a user accrues over years and rides every hardware refresh intact; recovery from device loss that restores not just files but the agent's accumulated behavioral state and

history; a verifiable, append-only disclosure record letting a user (or, where applicable, an auditor) see every off-device transmission the agent ever made; per-context separation so a professional corpus and a household corpus never bleed together under one identity; and a private personal corpus that, by default, never leaves devices the user holds. Federation lets a draft begin on a phone and continue on a laptop as the same agent's work, with application-mediated handoff propagating the artifact across devices within the same federated scope.

Boundary Conditions

U.S. Provisional Application No. 64/070,239 is a provisional, early-stage filing; it establishes priority for the disclosed architecture but is not an issued patent, and claim scope will be determined in prosecution of any non-provisional application claiming its benefit. The underlying building blocks are not claimed as new: parameter-efficient fine-tuning (low-rank adaptation, prefix and prompt tuning), on-device model execution, secure enclaves and trusted platform modules, and cryptographic hash chaining are established prior art. What is disclosed is the architecture that composes them, an agent that persists as the durable, identity-bearing object with models and devices subordinate to it.

Real constraints apply. On-device inference and local fine-tuning run within a device's bounded memory, compute, power, and thermal envelope, which the resource governance subsystem manages through budgets, scheduling, quiescence, and eviction; a wearable will carry a smaller, more quantized registry than a desktop. Hardware-anchored identity binding presumes a usable hardware security element on each device. Governed migration depends on attestation from both devices, so a destination that fails hardware, governance, or trust criteria can be refused, which is a security property and a usability limit at once. The strength of the privacy invariant depends on its enforcement mechanisms (an egress filter, per-component isolation, hardware-anchored attestation) being correctly implemented and not circumvented by a

compromised runtime. None of these are described here with performance figures, and none are asserted in this article; the application discloses mechanisms, not benchmarks.

Disclosure Scope

The architecture, mechanisms, and guarantees described here (hardware-anchored identity binding, governed migration, the federated agent identity record, the personal corpus model, the append-only lineage under a continuity proof, scope partitioning, and the privacy invariant) are disclosed in U.S. Provisional Application No. 64/070,239, "Agent-Resident Execution Substrate with Governed Inference Tool Registry and Lineage-Derived Personal Corpus Model Training," and in the sibling portfolio applications it incorporates by reference. The personal cross-device deployment scenario, the consumer device tiers (phone, laptop, tablet, wearable), and the market framing around device loss and upgrade are presented as a faithful enabling implementation of that disclosure and as external context; they are not themselves patent claims, and nothing here should be read as defining or limiting the claims of any application. References to hardware components such as secure enclaves and trusted platform modules, and to techniques such as parameter-efficient fine-tuning, denote established, externally defined technology used by the architecture, not inventions claimed by it.

Agent-Resident Execution Substrate (/agent-resident-execution-substrate)

All 40 steps → (/inventive-steps)

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)
- Heterogeneous Inference Endpoints (/articles/agent-resident-execution-substrate/heterogeneous-inference-endpoints)
- Atomic Lifecycle Substitution (/articles/agent-resident-execution-substrate/atomic-lifecycle-substitution)
- Integrity Signal Feedback (/articles/agent-resident-execution-substrate/integrity-signal-feedback)
- 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)
- Counterparty Identity Records (/articles/agent-resident-execution-substrate/counterparty-identity-records)
- 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)

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)**
- Enterprise Agent Fleets: Stable Agent Identity and Governed Tool Access Across Model Upgrades and Infrastructure Migration (/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)
- Edge and On-Device Agents: Hardware-Bound Identity Across Heterogeneous Inference Endpoints (/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)
- Governed Tool Lifecycles for Managed Inference-Provider Ecosystems: A Substrate Approach to Owning, Routing, and Retiring AI Tools (/articles/agent-resident-execution-substrate/managed-tool-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)

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