FINAL REPORT
MARCH 2026

# The Future of Personal AI

*Portable and Persistent Personal Memory through a Unified Human Context Protocol*

**Global Memory Workshop · March 11–12, 2026**
*Stanford Institute for Human-Centered Artificial Intelligence (HAI)*

# Executive Summary

Personalization increasingly shapes how digital systems support learning, health, and access to opportunity. Yet today's systems are largely platform-centric, fragmented across tools, and built on short-term behavioral signals that fail to represent authentic goals, constraints, and long-term intent.

This limitation becomes a binding constraint as generative AI is deployed in high-stakes environments: despite large language models' exceptional reasoning capabilities, they can only provide meaningful support when grounded in accurate, current, and appropriately scoped personal context. At the same time, the locus of value in digital systems is shifting — durable advantage increasingly comes not from network scale alone, but from sustained understanding of individuals over time.

Without intervention, this shift may consolidate durable personal context within proprietary "cognitive layers," with long-term implications for portability, interoperability, and user agency. Given the alarmingly rapid accumulation of personal contexts in AI service, how to create an ecosystem that facilitates more user agency and control over their personal data becomes an urgent question.

> **CORE THESIS**
>
> *Advancing the next generation of personalization requires a new foundation: portable, persistent personal memory — a user-governed layer of context that persists across interactions, moves across tools, and can be shared selectively by role, consent, and purpose.*

This report synthesizes the pre-read materials and frameworks developed for the Global Memory Workshop co-organized by the Gates Foundation, Mozilla Foundation, Stanford HAI and AWS, presenting an analysis of the current personalization landscape, a concrete architecture for user-governed memory, key applications in education and health, and a priority agenda for standards, evaluation, infrastructure, and policy.

# 1. The Personalization Status Quo

For the past three decades, more and more of what makes us who we are has been captured and digitized. Travel patterns live with Uber and Expedia. Purchase histories sit inside Amazon. Meta owns social networks; Google holds search intent. Each platform has built a remarkable personalization engine on a narrow slice of human context — and each slice is locked inside the platform that captured it.

From the modeling perspective, every domain has had its own specialized engine interpreting its own specialized data: collaborative filtering for movies, ranking models for search, graph algorithms for social networks. Context and interpretation were tightly coupled, and the platform that collected the data was the only one that could make sense of it. The result: a digital life that is siloed, non-portable, and incapable of representing who we are across the contexts where we actually live.

## 1.1 Four Structural Limitations

- User contexts are largely siloed. Each system builds and maintains its own partial representation, with limited ability for context to persist or transfer across tools, programs, or environments.
- Current user models fail to capture dynamics and long-term intent. Systems are effective at modeling recent behavior but fail to maintain a longitudinal profile that captures nuanced goals. Many real-

world objectives — learning a skill, managing a chronic condition, supporting a family member — are ongoing projects requiring continuity.

- Incentive structures skew personalization toward short-term engagement. When engagement metrics serve as the primary optimization target, systems learn what elicits reliable reactions instead of what supports sustained progress.
- Most personalization systems provide limited user agency. Individuals have little visibility into how they are represented, limited ability to correct those representations, and few practical controls that travel across platforms.

### 1.2 Why This Matters Now

Two things are changing fundamentally at the same time. First, AI systems are becoming the ultimate context observers. People share things with a chatbot they would never type into a search box: health issues, career doubts, relationship dynamics, half-formed goals. A single conversation can contain more nuanced personal context than years of click logs.

Second, large language models are becoming general-purpose context interpreters. A general-purpose model can now reason over calendars, emails, medical notes, learning histories, and purchase records in natural language. Interpretation is sometimes less precise than a purpose-built model, but dramatically more flexible and far easier to deploy.

Personalized services have always been the product of personal context times a reasoning algorithm. For most of the internet era, the algorithm was the scarce and defensible asset. In the agentic era, reasoning is increasingly abundant — and context becomes the binding constraint. Whoever holds the richest, most current, most trustworthy picture of a person will deliver the most valuable personalized services.

## 2. The Vision: Portable, Persistent Personal Memory

We consider an alternative model for personalization in which personal context is governed by the individual rather than embedded within platform-specific stacks. In this model, a single evolving memory layer persists across tools and interactions and is shared safely and selectively. The core shift is from platform-centric personalization — where preference representations remain internal and non-portable — to user-controlled continuity, in which access is narrowly scoped, purpose-limited, and revocable.

### 2.1 Prior Efforts and Persistent Obstacles

The concept of unified user memory is not new. Prior efforts — including personal data stores, pods, and related mechanisms — have sought to give individuals control over storage and permissions. However, these approaches have faced substantial adoption barriers and ecosystem constraints. For instance, solid-style pods provide user-controlled decentralized storage with revocable permissions but have been constrained by self-hosting frictions, an immature ecosystem, and limited support for natural-language scoping.

Across these efforts, three product and systems challenges recur:

- Low-friction context contribution: users require a mechanism to contribute context naturally; systems that rely on deliberate, ongoing data ingestion impose a usability tax that does not scale.
- Lifecycle management: a persistent memory substrate necessarily contains both stable constraints and temporally contingent preferences, and without lifecycle operations (update, deprecate, correct), memory becomes stale or erroneous.

- Safe and useful routing: context must be routed across heterogeneous services in a way that is both useful and safe — requiring minimization, authorization enforcement, and purpose limitation.

## 2.2 Why an LLM-Native Approach Changes Feasibility

Two developments make a user-governed memory substrate more plausible now than earlier paradigms. First, natural language has become a primary interface modality for AI systems, reducing the cost of expressing and updating preferences. Second, agentic workflows increasingly require cross-application execution, raising the value of a central context layer that reduces repeated onboarding and supports persistent task state.

Together, these trends favor memory representations and interfaces that are: (a) human-readable and editable, (b) readily consumable by language models, and (c) enforceable through scoped access controls at query time.

## 2.3 Architecture: A User-Governed Context Layer

We model the memory layer as a control plane that intermediates between user context and consuming applications. The essential property is mediated access: multiple assistants and applications query a shared store through a policy-enforcing control layer that returns only authorized, minimal, task-relevant context.

| Component | Function |
|---|---|
| Memory Store | Durable store holding human-readable entries augmented with machine-usable metadata (timestamp, source, category, confidence) and optional vector index for semantic retrieval. |
| Policy & Authorization Service | Validates tokens and enforces least-privilege scopes. Category-scoped grants with revocation and differentiation between temporary and persistent sharing. |
| Context Router (Memory Manager) | Selects what to disclose given current request, permitted scopes, and minimization rules. Enforces minimum necessary disclosure by selecting relevant items and redacting unrelated fields at inference time. |
| Audit & Provenance Subsystem | Append-only log of memory accesses and mutations, enabling user inspection of which agents accessed which data and when. |
| Interoperability Interface | Tool interface enabling assistants and third-party applications to perform constrained memory operations (add, search, update, delete) with authentication and authorization mediated by the policy layer. |

## 2.4 Memory Flow: Observation, Digestion, and Serving

### Observation (Permissioned Capture)

The system captures explicit preferences, constraints, accessibility needs, communication styles, long-term goals, and project context under explicit user permission. Observations are recorded as candidate memory items with provenance so they can be inspected and corrected.

### Digestion (Structuring and Maintenance)

Observation produces raw signals; digestion transforms them into durable and queryable memory. This includes entity linking, redundancy reduction, summarization, and conflict handling. The design must

preserve a strict distinction between user-declared facts and model inferences, and keep inferences visible and correctable to mitigate identity hardening from transient behavior.

**Serving (Minimum Necessary Context)**

Serving is the principal trust surface. Systems receive only the context required for a defined task, purpose, and duration, rather than bulk export of a profile. Fine-grained, revocable permissions enable role-appropriate sharing and allow users to revoke access at any time.

### 2.5 Integration with Third-Party Applications

A user-governed memory layer must support routine use by third-party applications through a standardized interaction pattern covering login, authorization, context retrieval, and context updates. Two invariants are non-negotiable: third-party applications never receive unconstrained access — retrieval is always conditioned on authorization and minimization at the time of use; and the system's trustworthiness depends on enforceable, auditable operations rather than voluntary downstream compliance.

## 3. Applications

Portable personal memory shifts personalization from optimizing engagement to supporting sustained progress over time. Its value is greatest in domains where outcomes depend on continuity of context and where fragmentation imposes real costs — most notably education and health.

### 3.1 Education: Enabling Learning Continuity Across Systems

Education is a long-horizon process, but most learning technologies are optimized for short-term interactions within individual platforms. Learner context — goals, prior knowledge, constraints, and progress — does not persist reliably across time or tools. As a result, personalization resets at every transition, instructional coherence breaks down, and learning systems struggle to support cumulative development. This fragmentation limits effectiveness and increases inequity, particularly for learners who move between programs, institutions, or modalities.

A portable memory layer addresses this structural gap by separating learner context from individual applications and placing it under learner and family direction. If realized well, such a system would operate across institutional and geographic boundaries, reducing discontinuity at moments of transition and aligning personalization with long-term learner growth rather than short-term engagement.

> **USER STORY: ALICE'S PORTABLE LEARNER PROFILE**
>
> *Alice has been learning mathematics on 'Math Kingdom' for several months. When she decides to switch to 'Fantasy Academy,' instead of starting from scratch she authorizes scoped access to her portable learner profile — limited to math-related mastery signals and learning preferences. Within minutes, Fantasy Academy imports this context, maps it to its own curriculum, and begins personalized recommendations on day one. As Alice learns, new progress updates flow back into her learner profile — keeping her record consistent across platforms without either platform owning it.*

| Field | Alice's Portable Math Learner Profile |
|---|---|
| Goal / Timeline | Placement exam in 8 weeks |

| Field | Alice's Portable Math Learner Profile |
|---|---|
| Current Level | Precalculus; starting functions + trig review |
| Mastery Snapshot | Linear eq. 0.92; Factoring 0.78; Rational expr. 0.55; Word problems 0.48; Domain/range 0.60 |
| Error Patterns | Negative-sign distribution; inverse vs reciprocal; unit mismatch in setup |
| Learning Preferences | Worked examples first; graphs help; progressive hints |
| Pacing / Constraints | 25–30 min weekdays; 60 min weekends; short videos + exercises |
| Accessibility | Large text / high contrast |
| Provenance / Recency | Last assessed 2026-02-02; confidence: high (mastery), medium (errors) |

### 3.2 Health: Supporting Continuity of Care and Self-Care

Health is the domain where continuity is most critical — and where its absence creates the greatest risk. Individuals are repeatedly asked to restate symptoms, medications, allergies, prior diagnoses, and what has or has not worked, often during periods of stress or illness. Despite digitization, health personalization remains fragmented, with context scattered across tools and institutions and continuity managed largely by the patient.

A portable memory layer offers a systems-level response — if designed with strict consent, compartmentalization, and purpose-bound access.

> **USE CASE: ETHAN'S VACCINATION RECORD**
>
> *When Ethan's family moves to a new state, his new pediatric clinic uses a different portal. Instead of requesting paper copies, Ethan's parent authorizes the new clinic to access a limited portion of his portable health record — scoped specifically to vaccination data. The clinic immediately imports the verified immunization history, determines which vaccines are complete and which are upcoming, and requires no duplicate shots or additional paperwork. Neither health system owns the record — they simply contribute to and retrieve information with consent.*

| Field | Ethan's Portable Vaccination Profile |
|---|---|
| Care Goal / Requirement | Meet state school immunization requirements |
| Current Status | Up to date through age 10; next Tdap due in 6 months |
| Vaccination Snapshot | MMR ✓; DTaP ✓; Varicella ✓; Hep B ✓; Flu (last dose 2026-10-12) |
| Pending / Upcoming | Tdap booster due 2026-08; HPV series not yet started |
| Consent Scope | Immunization verification with schools; full record with licensed providers only |
| Provenance / Recency | Last updated 2026-02-02; from licensed clinic EHR; status: confirmed |

## 4. Risks and Challenges

The same logic that makes portable context valuable also makes centralized, proprietary context dangerous. If the memory layer consolidates inside a handful of frontier model providers, switching costs will rise rather than fall. Public-interest systems — schools, clinics, social services — will find themselves renting access to their own users' context. And the people whose lives are encoded in these systems will have even less visibility into what is known about them than they do today.

## 4.1 Technical Challenges

The technical challenges are real but tractable. Memory systems need provenance on every entry — origin, confidence, modification history — so that downstream decisions can be audited. They need to degrade gracefully when retrieval is incomplete or conflicting. They need to run on accessible hardware, including low-bandwidth and on-device settings. They need longitudinal evaluation frameworks, because today's benchmarks measure isolated tasks while memory systems accumulate influence over months and years. And they need interoperable schemas, because portability without standards is just a promise.

## 4.2 Legal and Structural Challenges

The legal and structural challenges may be harder. Existing incentives push in the wrong direction: the platforms best positioned to build rich context layers are also the ones that benefit most from keeping context non-portable. Data portability rights under GDPR and CCPA point toward user-controlled context as a legal requirement, but the technical implementations that would make those rights meaningful do not yet exist at scale.

There is a deeper question about how human–AI interaction should be treated as a category. The intimacy of these exchanges warrants a new class of protection, similar in spirit to privileged professional communications. Human–AI interaction data is more concentrated, more revealing, and more actionable than any previous category of personal data — and the legal frameworks we have were not written with it in mind.

## 4.3 On-Device and Small Model Constraints

Most deployed memory systems currently rely on API-based models to perform summarization, filtering, retrieval orchestration, and validation. This introduces privacy concerns because sensitive user context must be transmitted to external services. Ideally, the memory manager should run close to the memory stored on-device or within user-controlled environments.

Key next steps include: (i) distillation and fine-tuning toward memory tasks; (ii) hybrid pipelines combining small models with deterministic components (rules, schema validators, policy filters); and (iii) tiered execution, where local models handle sensitive operations and only escalate to larger models under explicit permission.

## 4.4 Security, Privacy and Public Trust

Given the sensitivity of personal data, security and privacy remain the biggest concerns — both as a technical problem and a deep social challenge. To address this beyond technology, we need transparent governance frameworks that clearly communicate what data is collected, how it is stored, and who can access it. Public trust ultimately depends not only on the robustness of the underlying infrastructure, but on whether communities feel genuine agency over their own data. Building that trust will require sustained engagement with civil society organizations, local institutions, and the populations these systems are designed to serve.

# 5. Discussion: Priority Gaps

While a portable personal memory layer appears conceptually simple, achieving it requires coordinated progress across standards, evaluation, systems, and policy. The following summarizes the primary gaps identified by workshop participants.

## 5.1 Industry Standards for Portable Personal Memory

Any interoperability standard must distinguish between (a) portable representations of user-declared state, and (b) model-generated inferences. Current 'memory' implementations remain product-defined and non-interoperable. A minimal standard should specify:

- Common schemas for preferences, goals, constraints, relationships, and facts — with provenance, uncertainty, and temporal validity.
- Interoperable APIs for read/search/write/update/delete with versioning.
- A consistent permission model with scoped, revocable access and auditable consent.
- Export/import formats that preserve meaning rather than raw logs, with access and mutation audit trails.

**REFLECTION QUESTIONS**

- What is the minimal interoperable schema required to support a portable learner theory of mind — including goals, mastery state, preferences, constraints, provenance, confidence, and temporal validity — while preserving innovation and pedagogical flexibility?
- What should be the default unit of permissioning (domain-, purpose-, role-, or session-level), and how should consent withdrawal propagate across interconnected tools and institutions?
- How can portable memory give students and families more agency when the current state is for schools and districts to maintain ownership of student records?
- How does a targeted universalism approach — designing for mobility, vulnerability, and fragmentation — influence the development of infrastructure that is more interoperable, durable, and equitable?

## 5.2 Evaluation and Benchmarks

Memory-driven personalization is difficult to benchmark using standard offline datasets. Effective evaluation typically requires individualized, longitudinal feedback: whether retrieved context was relevant, whether it improved outcomes, whether it over-shared, and whether stored memory remained correct as user preferences evolved.

A practical benchmark suite should include: (i) task utility metrics; (ii) minimization metrics; (iii) stability metrics (memory drift, contradiction rate, error accumulation); (iv) repairability metrics; and (v) adversarial robustness (prompt injection, inference attacks, and data poisoning).

**REFLECTION QUESTIONS**

- What is the 'right' primary metric — utility gain, over-sharing rate, or user trust — and how should tradeoffs be reported across them?

- How can we design privacy-preserving longitudinal evaluation (e.g., on-device measurement, federated aggregates) without losing interpretability and reproducibility?
- How can portable memory grant granular permissions ('share my math grades but not English'); purpose-based controls; preview before sharing; and easy revocation?

### 5.3 Governance, Transparency, and Equity

Adoption depends on governance structures that make portability trustworthy, not merely possible. Existing frameworks such as FERPA, COPPA, and CIPA provide a baseline, but cross-platform memory systems will require additional enforceable safeguards: purpose limitation by design, granular and revocable consent, auditability, restrictions on secondary use, and clear allocation of liability across actors.

Equally critical is a shared accountability and correction framework. Because personal memory is longitudinal and interoperable, inaccuracies or misuse could propagate across systems. In high-stakes domains like education and health, portability must be paired with accountability to avoid shifting risk onto users.

**REFLECTION QUESTIONS**

- What governance mechanisms would most effectively ensure that personal memory systems are used in ways that align with user interests (e.g., purpose limitation, transparency requirements, auditability)?
- When personal memory is inaccurate, outdated, or used in unintended ways, what shared accountability and correction framework should be in place?
- What design and governance considerations are needed given that schools/districts are selecting learning tools and 'consenting' on behalf of students and families?
- How might parent and student voice meaningfully shape transparency, consent processes, and correction pathways in AI-enabled systems?
- How should accountability be defined in cases of data breaches or AI-related harm, particularly in ways that strengthen parent and student trust?

## 6. Next Steps: From Architecture to Infrastructure

Three things would move this from a promising architecture to working infrastructure.

### Standards First

Workshop participants were emphatic that open standards are a prerequisite rather than a downstream concern. Retrieval interfaces, consent representations, provenance metadata, and portability schemas all need open specifications that can evolve without breaking legacy clients. Academic and civil-society voices matter here precisely because incumbents have weak incentives to converge.

### Real Pilots in Public-Interest Domains

Education, health, and workforce systems are where the value of portable context is most legible and where the risks of getting it wrong are most serious. Pilots in these settings will do more to build trust and shape standards than any amount of whitepaper discussion.

### A Legal and Policy Foundation

Federal data and privacy policy was named by workshop participants as a prerequisite rather than an afterthought. The specific question of how to treat human–AI interaction data deserves its own consideration: it is more concentrated, more revealing, and more actionable than any previous category of personal data, and the legal frameworks we have were not written with it in mind.

For example, human–AI interaction deserves something closer to attorney–client privilege: it is uniquely intimate, uniquely concentrated, and currently almost entirely unprotected.

## Conclusion

The locus of advantage in AI is shifting from model capability alone to the ability to accumulate and responsibly apply longitudinal understanding of individuals. That shift is already underway, and it will happen with or without a user-governed layer. The question is whether the context that increasingly defines our digital lives will live inside infrastructure we control, or inside infrastructure that controls us.

There is a narrow window to build the former before the latter becomes the default, and the work to do it is concrete, collaborative, and ready to begin.

> **THE CORE CHOICE**
>
> *If personal context can be aggregated, governed, and served from a place the user actually controls, we unlock a new generation of services across learning, health, entertainment, and the public sector — services that adapt to real people over real time horizons, rather than starting from zero in every new app.*

*Note: This workshop and report were not designed to facilitate policy development or regulatory guidance. Policy and governance related questions naturally arose during the discussion and we hope this report will inspire broader discussions on AI policy and regulations related to personal data.*

Global Memory Workshop · March 11–12, 2026 · Gates Foundation · AWS · Mozilla Foundation · Slalom · Stanford HAI