

RPTL: The Rotating Privacy Translation Layer

A privacy architecture for Personal AI

Gloabi / MySALT AI Inc. | Public White Paper, Version 1.0 | Published August 11, 2026 | Patent pending | gloabi.com

Current scope: RPTL applies to supported text-based Personal AI interactions. Real-time voice-to-voice conversations powered by Gemini Native Live require audio to be processed by that voice service and are not currently covered by RPTL.

**A Personal AI should understand the user it serves.
The models powering it should not need to know who that
user is.**

Abstract. Personal AI is most useful when it understands a person's real world: the people they know, the plans they are making, and the information needed to complete a task. That same context can create a privacy problem when it is sent directly to a language model or an inference provider. RPTL, the Rotating Privacy Translation Layer, is Gloabi's privacy boundary between a Personal AI application and the models that generate responses. Before protected identifying references are sent for inference, RPTL replaces them with session-specific aliases. The model can reason coherently about those aliases during the conversation. When a response comes back, Gloabi restores the real meaning for the user; when the user has approved an action, the real values are restored only at the application boundary where

that action executes. The result is an AI that can reason about a user's world without the inference model receiving the protected identity behind it.

**Private user context → rotating aliases → model reasoning →
restored meaning**

1. The Problem: Useful AI Has a Privacy Problem

The current AI experience commonly asks people to reveal context directly to the system answering them. A user may share a relative's name, an address, an email address, a phone number, or details about their work and personal life. In many conventional systems, those details become part of the prompt delivered to the model provider.

Traditional redaction is not a satisfying answer. Removing every meaningful reference can make a request unusable: an AI cannot meaningfully help a user communicate with a person, recall a plan, or perform an approved task when the crucial context has been replaced with generic blanks. At the other extreme, passing the original values through provides utility but unnecessarily exposes identity to every model in the path.


A Personal AI makes this tension especially important. Unlike a one-off search query, a Personal AI is expected to be useful over time. It needs continuity. It may help across web, mobile, desktop, phone, and messaging experiences. Privacy must therefore be part of the architecture rather than an afterthought applied only to a single model or channel.

2. Design Principles

- **Identity should be protected without destroying utility.** The AI should retain the relationship between references in a conversation without receiving the real-world identity behind them.
- **Continuity should not become cross-conversation tracking.** A model needs a stable reference during one conversation, but that reference should not persist as a permanent identifier in later conversations.
- **The protection should not depend on one provider.** A privacy boundary should work whether the request is handled by a self-hosted model or an external provider.
- **The user should receive the real result.** Aliases are for the inference boundary; the user sees their own real context, and approved actions operate on the real values only where needed.
- **Claims should remain honest.** RPTL reduces identity exposure. It does not claim that semantic context cannot ever reveal information, and it does not remove the need for consent, secure systems, and user control.

3. Architecture

RPTL is an application-layer translation system. It sits between the Personal AI experience and the inference service used to respond. It has three conceptual responsibilities: identify protected references, create a session-specific representation for them, and restore the meaning when the response returns.






Figure 1. Conceptual RPTL message flow. Protected identity is translated before inference and restored after the model response.

A. Session-scoped aliases

When RPTL identifies a protected reference, it associates that reference with an alias for the active conversation. The same reference resolves to the same alias within that session, allowing the model to maintain coherent context. A fresh session uses a different alias. The model therefore has a useful conversational reference without receiving a durable real-world identifier.

B. Reversible application boundary

The translation is reversible within the trusted application boundary. A response that refers to an alias is restored before it is shown to the user. If an approved workflow needs a real value, restoration occurs at the point the application executes that workflow, not in the model prompt. This separation makes it possible for a model to help formulate or plan an action without being given the protected value itself.

C. Model-independent operation

RPTL is designed as a layer around inference rather than as a modification to a particular model. This matters for reliability and choice. If an application changes models or uses a fallback provider, the conversation's protected references can remain consistent because the privacy session belongs to the application, not the model provider.

4. From Private Context to Useful Response

Consider a user asking their Personal AI to help prepare a message to a person in their life. Without RPTL, the person's identity and contact details may travel directly to the inference provider. With RPTL, the model receives a stable alias that preserves the conversational role of that person. It can draft a useful response, reason about prior references in the same conversation, and request an approved action. The application restores the real value only after the model's work is complete.

This approach is different from one-way masking. A generic placeholder such as [NAME] can remove too much meaning and tends to produce poor language. A stable, natural-language alias lets a model continue to reason about an entity as an entity. It is also different from a fixed tokenization vault: a fixed token can itself become a persistent marker that links every conversation in which it appears.

RPTL is built around a practical compromise: protect the identifying reference while preserving the non-identifying meaning needed to help the user. That design preserves the usefulness people expect from a Personal AI while narrowing what an inference model needs to receive.

5. Threat Model and Boundaries

RPTL is designed for the risk that an inference provider, its logs, or a compromised inference environment can observe prompts and responses. The layer aims to ensure that protected identifiers are represented by session-specific aliases rather

than delivered in their original form. It also reduces the ability to use a repeated alias as a long-term link across a user's separate conversations.

RPTL does not promise the impossible. In normal operation, the model still receives the non-identifying content of a user's request. Exceptionally distinctive context can sometimes imply information even when names and direct identifiers are removed. RPTL is therefore one part of a broader privacy posture that includes secure application design, appropriate data retention, explicit user consent, and controls over what actions an AI can take.

This scope is intentional. A credible privacy architecture should make a precise claim it can defend: a model should not need a user's direct protected identity in order to help with the user's request.

Extending identity protection to more modalities is an active area of development. Gloabi aims to improve detection and privacy coverage without relying on raw user conversations as a generalized training resource. Future work includes applying privacy-preserving translation more broadly to voice experiences while maintaining the real-time quality users expect.

6. Why Rotation Matters

A privacy substitute that never changes can become another identifier. If the same replacement appears in every conversation, an observer can group those conversations together even without knowing the underlying name. That persistent correlation undermines the privacy benefit of tokenization.

RPTL changes the scope of consistency. It maintains consistency only for as long as the model needs it: the active conversation. When a new conversation begins, the protected reference receives a different alias. This gives a model enough continuity to be useful while reducing the value of accumulated aliases as a cross-conversation tracking mechanism.

7. Practical Benefits

- **Better user trust.** Users can receive useful help without being asked to routinely provide direct identity information to each model involved in a request.
- **Provider flexibility.** The application can maintain a consistent privacy boundary across a changing mix of self-hosted and external inference services.
- **Practical task support.** The model can assist with language, planning, and approved workflows while the application preserves control over protected values.
- **A foundation for Personal AI.** The same privacy principle supports text interactions today and is designed to extend across desktop, mobile, and future voice experiences.

8. Validation Approach

Gloabi validates RPTL as an engineering system rather than relying solely on a conceptual promise. The validation focus includes correct round-trip restoration, stable references within a conversation, rotation across independent conversations, preservation of utility in normal AI interactions, and continuity when a request moves between available inference providers. The system is also designed with operational controls so that privacy features can be monitored and managed responsibly.

The purpose of public validation is not to publish implementation internals. It is to establish the operational principle: protected identity is translated before inference, real meaning is restored for the user, and the Personal AI remains useful through the process.

9. The Future of Personal AI Requires a Better Privacy Boundary

The future of AI will not be defined only by larger models. It will be defined by whether people can trust AI with the context that makes it useful. A Personal AI should be able to understand a user's world, help them communicate, and assist with real tasks without requiring every inference model to learn the identities involved.

RPTL offers a practical direction: translate protected identity at the application boundary, keep references coherent for the duration of a conversation, rotate those references across conversations, and restore the real-world meaning only where it belongs. It is designed to let a Personal AI be personal without making a user's identity the price of usefulness.

About Gloabi

Gloabi is building a Personal AI platform in which every user can have an AI with its own identity, practical capabilities, and a relationship with one person. Gloabi's goal is to make Personal AI accessible to non-technical people while treating privacy, user control, and real-world usefulness as core product requirements.

RPTL — Rotating Privacy Translation Layer. © 2026 MySALT AI Inc. All rights reserved. Gloabi is a product of MySALT AI Inc. — gloabi.com. Patent pending — U.S. provisional patent application filed August 2026. This public white paper describes the RPTL architecture at a conceptual level and intentionally omits implementation-specific security and operational details.