A Voice-First Agentic AI Architecture for Veteran Navigation, Transition, and Personalized Support

A human-centered framework demonstrated through the ValorBuddy platform

Eugene Ezenwa Ebem

Co-Founder & Chief Technology Officer
Tagus Technologies LLC
Dallas-Fort Worth, Texas, USA

PREPRINT / WORKING PAPER • September 2026

Keywords: Agentic AI; Veterans; Large Language Models; Voice AI; Human-Centered AI; Digital Assistants; Responsible AI

SCHOLARLY RECORD

ORCID: 0009-0005-8470-4922
DOI: 10.5281/zenodo.25881294
Version 1.0 | September 2026

Independent technical preprint

Abstract

Veterans and transitioning service members often interact with a fragmented ecosystem of benefits information, health resources, employment services, appointments, community organizations, family-support programs, and civilian-transition resources. This paper proposes a voice-first agentic artificial intelligence architecture designed to reduce navigation burden through personalized, conversational assistance. The architecture is demonstrated through ValorBuddy, a digital platform developed by Tagus Technologies LLC for Veterans, service members, and their families. Unlike conventional question-answering chatbots, the proposed model follows a continuous assistance pattern: Know Me → Guide Me → Connect Me → Help Me Act → Grow With Me. The paper presents the architectural model, agent orchestration approach, consent-based personalization strategy, security and responsible-AI controls, and an evaluation framework. The system is positioned as a navigation and action layer that augments - rather than replaces - Veteran-service professionals, government systems, clinicians, attorneys, or emergency services.

Research positioning
This manuscript describes an architecture and evaluation framework. It does not claim clinical efficacy, VA endorsement, benefits-adjudication authority, or validated outcome improvements. Those claims require independent empirical evaluation.

1. Introduction

Transition from military service to civilian life can require a person to coordinate benefits, health care, appointments, education, employment, family responsibilities, documents, community resources, and everyday logistics. The challenge is therefore not simply the availability of resources; it is often a problem of navigation, personalization, continuity, and action.

Advances in large language models and agentic AI create an opportunity to rethink the traditional digital-service interface. Rather than requiring users to repeatedly search websites, interpret documentation, and manually coordinate follow-up actions, an AI system can maintain authorized context and guide a user through multi-step tasks while preserving human oversight.

2. The Veteran Navigation Problem

A conventional digital interaction is transactional: User → Search → Information → User interprets → User acts. This places most cognitive and administrative responsibility on the Veteran.

A more assistive model is: User → AI understands authorized context → AI identifies need → AI guides → AI connects trusted resources → AI assists with action → AI retains only approved continuity. The objective is not merely to produce better answers, but to reduce the distance between a Veteran's need and a useful outcome.

3. From Conversational AI to Agentic AI

The proposed model is summarized as Know Me → Guide Me → Connect Me → Help Me Act → Grow With Me.

Know Me: build an authorized contextual understanding from information the user chooses to provide. Guide Me: interpret the objective and translate complex information into understandable next steps. Connect Me: route the user to authoritative or appropriate community resources. Help Me Act: support reminders, document preparation, task lists, resumes and other permitted workflows, with confirmation for consequential actions. Grow With Me: provide consent-based continuity while allowing users to review, correct or remove stored information.

4. Proposed System Architecture

The architecture is organized as a layered system in which interaction, reasoning, memory, actions, and governance remain separable and auditable.

1. Interaction Layer	Voice, text, accessibility, speech recognition and speech output.
2. Orchestration Layer	Intent/context routing, workflow decomposition and specialized-agent selection.
3. Intelligence Layer	LLM reasoning, retrieval, summarization and task-specific models.
4. Memory Layer	Session context, user-approved preferences, reminders and consent-managed continuity.
5. Action & Integration Layer	Resource discovery, reminders, documents, employment/transition workflows and authorized external actions.
6. Security & Governance Layer	Identity, least privilege, auditability, privacy, grounding, confirmation and human escalation.

Conceptual flow: Veteran / Family Member → Voice or Text Interface → Intent & Context Router → Specialized Agent(s) → Safety / Policy Validation → Response or Authorized Action → Auditable Outcome.

5. Agent Orchestration

A single general-purpose model does not need to perform every task. An orchestration layer can route requests to specialized capabilities such as benefits navigation, reminders, transition support, employment, wellness-resource navigation, travel/resource discovery, document assistance, and general conversation. The router determines whether a request requires retrieval, reasoning, memory access, an external tool, user confirmation, or human escalation.

6. Personalization and Memory

Persistent context can improve continuity, but it is also sensitive. The architecture distinguishes session context, user-approved long-term preferences, operational data such as reminders, and sensitive information requiring stronger controls. The governing principle is: remember what is useful, with permission, and protect what is sensitive. Users should understand what is retained and have meaningful control over it.

7. Security, Privacy, and Responsible AI

A Veteran-support AI platform may encounter sensitive personal information. Security must therefore be designed into identity, application, data, model, and infrastructure layers. Relevant controls include encryption in transit and at rest, least-privilege authorization, role-based access, secure secrets management, audit logging, session protection, data minimization, consent management, authoritative-source grounding, explicit confirmation for high-impact actions, and safe failure when confidence or authorization is insufficient. NIST's AI Risk Management Framework provides a useful structure for governing validity, reliability, safety, security, resilience, accountability, transparency, explainability, privacy, and harmful-bias risks.

8. Human-Centered AI

The architecture is not intended to replace Veteran Service Officers, social workers, clinicians, benefits specialists, caregivers, government personnel, attorneys, or community organizations. The AI functions as a navigation and coordination layer. Medical, mental-health, legal, financial, benefits-adjudication, emergency, and other high-impact situations require appropriate professional or emergency escalation.

9. ValorBuddy as an Implementation Context

ValorBuddy provides an implementation context for the architecture. The platform is designed around a voice-first conversational experience for Veterans, service members, and families, with capabilities spanning AI conversation, reminders, personalized check-ins, resource navigation, employment and transition assistance, document organization, travel support, and other daily-life workflows. The research value of the platform depends on measurable reductions in navigation effort and improvements in safe task completion rather than on feature count alone.

10. Evaluation Framework

A credible evaluation should measure task completion rate, navigation reduction, successful resource connection, action completion, response groundedness, escalation accuracy, user trust and satisfaction, and safety error rate. Task Completion Rate can be defined as successfully completed tasks divided by total attempted tasks. Navigation Reduction Rate can compare interaction steps under conventional navigation with agent-assisted navigation. Safety evaluation should separately track materially incorrect, unsafe, unauthorized, or misleading outputs.

11. Research Questions

RQ1: Does voice-first agentic assistance reduce the steps required to locate appropriate Veteran resources?

RQ2: Does consent-based personalization improve task completion compared with stateless conversational systems?

RQ3: Which Veteran-support tasks benefit most from agentic assistance?

RQ4: What degree of autonomy best balances convenience, trust, privacy, and safety?

RQ5: How effectively can authoritative-source grounding reduce hallucination and misinformation?

12. Limitations

Large language models can produce incorrect or misleading information, and agentic systems introduce additional security risk because they may perform actions rather than merely generate text. Personalization creates privacy obligations. This architecture does not itself demonstrate improved health, benefits, employment, or transition outcomes; those claims require controlled empirical evaluation. Evaluation should begin with lower-risk navigation and productivity tasks before expansion into higher-impact domains.

13. Future Work

Future research should include structured pilot studies, retrieval-augmented generation from authoritative sources, accessibility optimization, multilingual interaction, explainable agent decisions, interoperability with authorized systems, and quantitative comparison of conventional web navigation with agent-assisted workflows. A Veteran Navigation Benchmark containing representative multi-step tasks could enable repeatable assessment of accuracy, safety, efficiency, groundedness, and completion.

14. Conclusion

Voice-first agentic AI offers a model in which natural conversation, contextual reasoning, consent-based personalization, trusted retrieval, specialized agents, and controlled actions can provide continuity across a Veteran's journey. The central research question is not whether AI can answer Veterans' questions; it is whether responsibly designed AI can reduce the distance between a Veteran's need and meaningful action while preserving privacy, security, human oversight, and user agency.

References

[1] National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1, 2023.

[2] National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile, NIST AI 600-1, 2024.

[3] U.S. Department of Veterans Affairs. VA Artificial Intelligence. https://department.va.gov/ai/

[4] U.S. Department of Veterans Affairs. Whole Health. https://www.va.gov/wholehealth/

Author Note

Eugene Ezenwa Ebem is Co-Founder and Chief Technology Officer of Tagus Technologies LLC. This manuscript is an independent technical preprint describing a proposed architecture and research agenda. References to ValorBuddy identify the implementation context and do not imply endorsement by the U.S. Department of Veterans Affairs or any other government agency.