Security Whitepaper

Protected B–Level Safeguards for Voice Agent Platform

Version 1.1

Prepared for Clinical and Service Organizations in Canada

Executive Summary

This whitepaper outlines the security architecture, data handling practices, and operational Safeguards that support Protected B–level requirements for organizations using our AI voice agent platform.

The system is engineered to protect sensitive and health related information handled during calls, intake, scheduling, and operational workflows.

The purpose of this document is to provide a transparent overview of how confidentiality, integrity, and availability are maintained throughout the full lifecycle of voice related data.

1. Security Principles

Our platform is designed around the three core pillars required for handling sensitive Canadian data:

- Confidentiality – ensuring all data remains private, encrypted, and contained within approved environments.
- Integrity – ensuring data is accurate, tamper resistant, and handled using strong authentication and access controls.
- Availability – ensuring services remain reliable, redundant, and recoverable at all times.

2. Data Residency & Confidentiality Controls

- Canadian Data Residency: All persistent data, including call summaries, operational metadata, and audio artifacts, remains exclusively in the Canada (Central) AWS Region.
- Isolated Network Design: All core components operate inside an isolated Virtual Private Cloud (VPC) with no public exposure of databases or internal services.
- Private Connectivity: Communication with AWS services such as storage or transcription Endpoints is performed through VPC Endpoints (Private Link), preventing data flow through public networks.

3. Encryption & Integrity Controls

- Encryption at Rest: All storage systems including database, caching layer, and file storage are encrypted using organization controlled KMS keys.
- Encryption in Transit: All communication channels enforce TLS/SSL to protect data during use.
- Role Based Access Control: System access adheres to least privilege principles using strict IAM policies, ensuring only authorized components and service roles may access sensitive data.
- Audit Logging: Operational logs are captured, encrypted, and stored in accordance with Protected B expectations to ensure traceability and enforce accountability.

4. Availability, Redundancy & Reliability

- Multi AZ Deployment: Core services including the PostgreSQL database and Redis caching Layer operate in Multi Availability Zone mode to ensure continuity during infrastructure disruptions.
- Automated Backups & Recovery: Regular backups preserve data integrity and enable rapid restoration.
- Scalable Compute Layer: Application nodes run behind an Application Load Balancer (ALB) and Auto Scaling Group (ASG), ensuring stable performance during high call volume or operational bursts.

5. Operational Safeguards

- Continuous Monitoring: System health, error rates, and performance metrics are monitored in real time.
- Incident Response Alignment: Operational practices support rapid mitigation of issues affecting confidentiality, integrity, or availability.
- Segregation of Duties: Administrative access is separated from system level operations to minimize risk.

6. Summary & Assurance

The voice agent platform applies a layered security model that aligns with Protected B–level safeguards.

Through strict residency controls, encryption, isolation, redundancy, and governance practices, the architecture supports the secure handling of sensitive call related information in Canadian clinical and service environments.

This whitepaper may be shared with compliance teams, IT administrators, and operational leaders as part of procurement, onboarding, or due diligence processes.