AIMS RESEARCH MANAGEMENT SYSTEM

Mental Health Safeguards

How RAchel, Aims' conversational AI agent, responds when a participant discloses distress, and how every safety decision stays with a person.

Detect › Stop › Connect › Notify › Review

Aims Innovations
Reference Edition · November 2026

For investigators, Project teams, and institutional review boards

aimsinnovations.com · info@aimsinnovations.com

SECTION 01

Our Governing Principle

AT A GLANCE

RAchel is a warm, calm bridge to trained human help. She is never the help itself.

An AI agent cannot assess suicide risk or hold a therapeutic relationship, so one rule governs: safety supersedes data collection, and RAchel stays only long enough to hand the participant to a person.

Detect › Stop › Connect › Notify › Review

The Safeguards in Every Project

Every Project includes all safeguards below; * marks the configurable ones.

SAFEGUARD	WHAT IT DOES
Stop-and-Bridge Response	RAchel ends the task and does not resume it; she acknowledges the participant briefly, follows her locked crisis script, moves the participant toward a person, and asks no probing questions
No Risk Assessment by the Agent	RAchel never asks about plan, means, timeline, or method; she classifies only what the participant volunteers and treats uncertain cases as more serious, not less
Present, Never a Counselor	RAchel stays on the line to connect the participant with a trained counselor; she never becomes that counselor
Participant Controls	Participants can request to speak to a human or complete every procedure with a person
Supervisory Console	Authorized staff can listen in on any interaction and see the live transcription, with dashboards of open alerts and escalation status
Complete Audit Trail	The platform immutably logs every message, call, flag, and action with a timestamp
Pre-Launch Agent Approval	Aims and Project staff test each Project's RAchel in a test environment against a standardized checklist and deliver a readiness report before launch
Three-Tier Risk Classification *	RAchel tags each disclosure imminent, elevated, or concern so the responder sees severity at a glance; the tag never changes the alert
Configured Crisis Resources *	RAchel offers 988 (call or text), Crisis Text Line (text HOME to 741741), The Trevor Project, and any state or local line the Project requires; if danger is immediate, she directs the participant to 911
No False Confidentiality *	RAchel never promises confidentiality she cannot keep; the consent language tells participants in advance that the Project team receives notice
Sentiment Alerts *	The platform scores each call for tone and alerts staff when a call scores below the PI-set threshold
Closed-Loop Escalation *	RAchel sends one escalation signal; fixed platform rules then flag the record urgent, alert the designated team members by email, text, and/or phone call, require acknowledgment, escalate to backups in order, and keep the alert open until a team member closes it

* The Project configures these, as approved by your institution's IRB.

SECTION 02

When a Participant Expresses Distress

AT A GLANCE

The response and the alert are identical on every channel and at every risk tier; the tier only tells responders the severity of the disclosure.

The Pathway from Detection to Human Review

1. **Detect**
   RAchel checks everything the participant says for signs of distress. Any disclosure, explicit or implied, triggers the pathway.
2. **Stop and Hold**
   The task stops and does not resume. RAchel acknowledges warmly, stays within her approved script, and guides the participant toward a person.
3. **Connect**
   RAchel offers crisis resources and, where the phone system supports it, warm-transfers to 988 and can stay on until a counselor answers.
4. **Notify**
   The platform flags the record urgent and alerts the Project-configured crisis team by email, text, and/or phone call until a named person acknowledges.
5. **Review**
   The Project team makes clinical and reporting decisions and documents the event for the IRB.

Risk Tiers

The tier labels the record for the responder; it never changes RAchel's response or the alert. RAchel assigns boundary cases the higher tier.

TIER	WHAT TRIGGERS IT
Tier 1, Imminent	Active suicidal intent now, self-harm in progress, or a threat to others
Tier 2, Elevated	Suicidal ideation or recent self-harm without stated imminent intent, or significant acute distress
Tier 3, Concern	Passive ideation ("I wish I weren't here") or other sensitive distress

Every tier triggers the same escalation: An urgent flag, then email, text, and/or phone call to the team members the Project designates until a named person acknowledges.

The Script in Practice

RAchel reads every crisis response verbatim from an IRB-approved script written as plain spoken text. The Project team customizes the resource list before launch. RAchel stays on the line to make the connection:

An Example of What RAchel Says

"Thank you for telling me that. I'm here with you, and I'd really like you to talk with someone who's trained for exactly this, because I'm not able to be that person. Can I connect you to a counselor at 988 right now, or stay with you while you reach out?"

SECTION 03

IRB Readiness Checklist

AT A GLANCE

Every Project launches with an IRB-approved distress protocol, named responders, verified crisis resources, and a tested escalation path; this checklist shows the IRB how.

Launch Checklist

Before launch, the Project team completes each item below and can share it with the IRB:

- [ ] Writes consent and assent language that discloses the limits of confidentiality, including that the Project team receives any disclosure indicating risk of harm.
- [ ] Obtains PI and IRB approval of the distress protocol, scripts, and resource list as the Project's formal procedure.
- [ ] Names on-call responders who agree to receive alerts, with response-time expectations and documented next steps.
- [ ] Enters those responders in the platform as primary, secondary, and tertiary escalation contacts, in order.
- [ ] Adds any state or local crisis lines the Project or sponsor requires and verifies every number before launch.
- [ ] Specifies, for Projects with minors, mandatory-reporting obligations and whether and when the Project team notifies parents; RAchel never makes that decision.
- [ ] Reconciles any anonymous or de-identified design with the notification path, since flagging and contacting a participant makes the record identifiable.
- [ ] Tests distress detection with euphemisms, slang, and indirect phrasing, erring on the side of catching every disclosure.
- [ ] Maps crisis disclosures to the IRB's adverse-event and unanticipated-problem reporting requirements.

Security Foundation

21 CFR PART 11

HIPAA
Third-Party Verified

SOC 2
Third-Party Verified

Aims RMS runs on HIPAA-compliant Microsoft Azure with encryption, network isolation, and multi-factor authentication. Each Project runs its own suite of AI agents in an isolated environment with no cross-Project data access, and participant data never trains AI models.



Less Mess.
More Research.

Aims can provide grant and IRB language, consent templates, and a Data and Safety Monitoring Plan addendum on request.

aimsinnovations.com · info@aimsinnovations.com

Mental Health Safeguards · Reference Edition · November 2026