

Questions your clients will ask for voice AI agent approval

Executive Summary

Prodigal has taken voice AI agents through approval with the top three debt buyers and leading national lenders. We know what compliance, risk, security, and vendor management teams need to see before a voice AI agent can go live.

Most often, AI approvals stall when a client cannot get comfortable with three things: **what the AI will do**, **what could go wrong**, and how the agency will **remain in control** after launch.

For every question, give the client **the boundary, the control, and the proof**. State what the agent can do, what prevents the wrong action, and how the client can verify it.

The questions your clients are likely to ask

Scope and control

What exact jobs will the AI handle, and what remains out of bounds?

How to answer: Define the approved use cases and exclusions with talk tracks. Include the portfolios, stages, channels, hours, or jurisdictions covered, plus the gate for expanding scope.

What they are looking for:

A fixed boundary that they can approve and enforce.

Which decisions are rule-based, which involve AI, and which require a human?

How to answer: Separate interpretation from action. Explain what the model may understand, what deterministic rules decide, and where a person takes over.

What they are looking for:

No hidden autonomy or unclear ownership.

What conditions force a transfer, stop, escalation, or manual review?

How to answer: Name the exact triggers, the action taken, and the safe fallback when a human is unavailable.

What they are looking for:

Predictable behavior when the conversation becomes risky.

Consumer and compliance

How does the agent identify itself and complete every required disclosure?

How to answer: State the exact point at which the disclosure is delivered, how completion is recorded, and what happens if the consumer interrupts or asks a question mid-disclosure.

What they are looking for:

Proof that the required language cannot be skipped.

How does it handle disputes, attorneys, bankruptcy, hardship, cease-and-desist, and transfer requests?

How to answer: Describe how each sensitive situation is detected, what the agent is prohibited from doing next, and the route for transfer or specialist review.

What they are looking for:

Compliance with the exceptions, not only the happy path.

What prevents an incorrect promise, offer, payment plan, or customer action?

How to answer: Identify the approved source of terms and the control that validates the action before it is communicated or executed. Make clear that the model cannot invent terms.

What they are looking for: A hard control between conversation and financial action.

What clients need to see before they approve

These questions test whether the system is safe in production and accountable after launch.

Data and model safety

Where does our data go, who can access it, and is any of it used to train models?

How to answer: Give a simple data-flow diagram covering systems, subprocessors, access roles, and retention.

What they are looking for:

No ambiguity about exposure, ownership, or reuse.

How is each client, portfolio, and customer interaction kept isolated?

How to answer: Explain the technical and operational separation: tenant boundaries, credentials, storage, session context, and deletion. Clarify how cross-client contamination is prevented.

What they are looking for:

Evidence that one client's data cannot affect another.

What happens when a model, telephony provider, CRM, or payment system fails?

How to answer: Describe the safe behavior for each dependency: what stops, what can continue, what the consumer hears, and when the call transfers or ends without taking action.

What they are looking for:

A safe failure mode, not a promise of perfect uptime.

Evidence and accountability

What testing proves the agent works in real conversations, not only controlled demos?

How to answer: Reference production-identical configurations, difficult scenarios, acceptance thresholds, and independent or manual review.

What they are looking for:

Evidence that covers variability and edge cases.

Can we inspect every call, decision, system action, and outcome?

How to answer: Show what is logged and reviewable: transcript, recording, identity state, disclosures, tool calls, rules applied, transfers, payments, and final disposition.

What they are looking for:

A complete audit trail, not a summary dashboard.

What changes require retesting, client notice, renewed approval, or rollback?

How to answer: Define the change-control categories and approval path. Include model, prompt, workflow, disclosure, integration, and business-rule changes, with a tested rollback process.

What they are looking for:

Control after launch, not only before it.

Client approval should not stall your AI rollout

We are happy to share the questions, evidence, and documentation that have helped move those reviews forward.

When approval becomes the bottleneck, we're ready to help.