

The Voice AI Decision Tree

Three decisions, two defensible defaults each, mapped to deployment speed and ROI. For most enterprise buyers, the faster default in each row is also the safer one.

Voice is the highest-volume, highest-cost, and most technically complex channel in customer service, and the one where AI adoption lags furthest behind. That gap is closing fast, and the organizations that move early hold a cost advantage that compounds quarter over quarter. The three barriers that made voice hard (latency, accuracy, and compliance) are now engineering problems with answers.

The decision tree

DECISION	FASTER DEFAULT (USUALLY ALSO SAFER)	SLOWER PATH
Telephony	Integrate with the existing carrier and stack.	Migrate carriers. Rarely worth it first.
IVR	Overlay and augment the existing IVR. Lower risk.	Rip and replace the IVR. Higher ceiling, longer path.
Escalation	Warm handoff that carries full context to the agent.	Cold handoff with no context. Avoid.

Why the faster default is usually the safer one

In the deployments Maven has measured, well-built voice agents resolve in the same range as chat on comparable intents. The complexity is real, but it's bounded. A short set of decisions, not the technology, now sets deployment speed. This framework pairs with the [90-Day Pilot Scope](#) for teams considering voice as their first automated channel.

The three barriers, and where they stand in 2026

BARRIER	WHAT IT MEANT ON A LIVE CALL	STATUS IN 2026
Latency	Sub-300ms response is non-negotiable on a live call.	Solved in production
Accuracy	Recognition across accents and overlapping speech.	Solved
Compliance	Requirements stricter than text; PCI on live calls.	Solved when built for it



What made voice hard in 2023 was latency, accuracy, and compliance. Those are now engineering problems with answers.