Short answer: sometimes—but not just because you can make the clone, or because the person once agreed to record a few lines. Using a real person's voice for a brand mascot can raise questions about consent, publicity and privacy rights, contracts, advertising rules, and AI-specific laws. The answer depends on who the person is, where the voice is used, what they agreed to, and how the synthetic voice is presented.

This is general information, not legal advice. Before launching a voice clone in advertising, customer support, or an app, get advice from a lawyer familiar with the relevant jurisdictions and media rights. The practical starting point is straightforward: obtain explicit, well-scoped **voice consent**, document it, and give the speaker a real say in how the clone is used.

Why brands are considering voice cloning

A mascot can make a brand more recognizable, and a consistent voice can carry that identity across ads, videos, games, and support experiences. A synthetic voice may also help a team produce new lines without scheduling another studio session each time. But convenience does not settle the rights question. A recognizable voice can have commercial value, and using it without permission may create legal and reputational problems even if the audio is technically generated from a model.

Research firms including Gartner and Grand View Research track the wider growth of AI, conversational systems, and voice technologies. That market attention helps explain why more companies are exploring voice agents and synthetic speech. It is not, however, a license to use any particular person's voice. A product trend and an individual's permission are separate things.

What "legal" depends on

There is no single global rule that says every voice clone is either allowed or prohibited. Laws vary by country and, in some places, by state or province. The same project may touch several legal areas at once. A lawyer should assess the full use—not just the model or recording process.

Publicity and personality rights

Many jurisdictions recognize some form of right to control commercial use of a person's identity. That identity may include a name, image, likeness, or other recognizable traits; whether and how a voice is covered varies. If consumers could reasonably identify the speaker, using a clone to sell products or represent a brand may require permission even when the person never recorded the final generated lines.

Rules can differ for living and deceased people, and rights after death vary considerably. Do not assume that a voice is free to use because the person has died, is famous, or cannot personally object. An estate, representative, or other rights holder may have a claim.

Contracts and the original recording

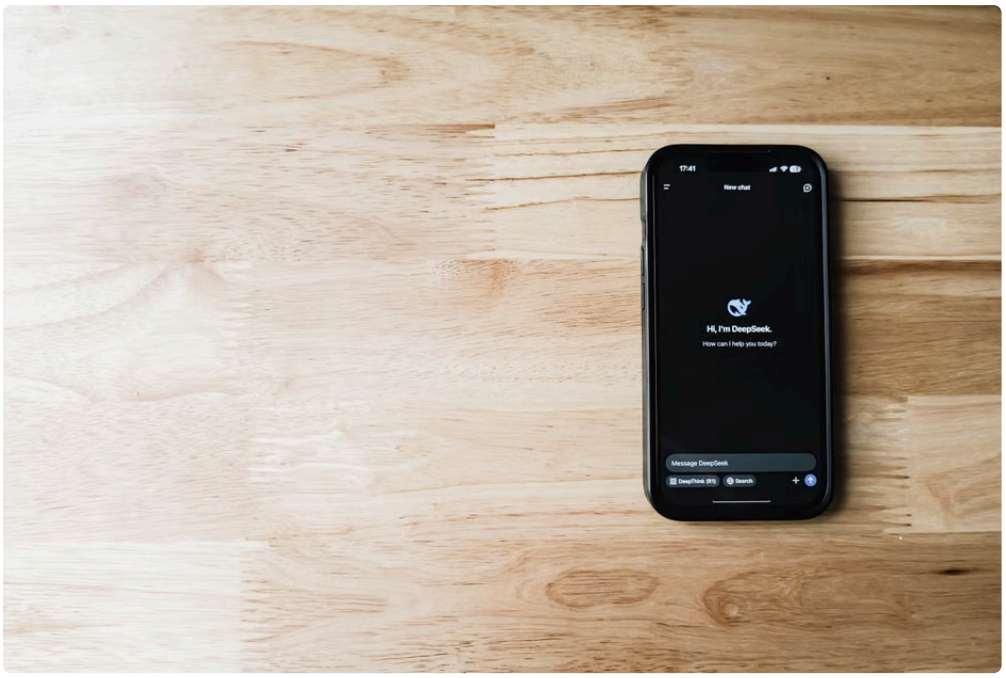
Permission to record a voice is not automatically permission to build a reusable voice model. A performer's agreement might cover a specific commercial, a set number of lines, or a defined campaign period—not a synthetic replica that can generate unlimited new dialogue. Review the recording agreement, talent release, union or guild requirements, and any rules from the production platform.

Contracts should make clear whether model training is allowed, what content the voice may produce, where it may appear, how long the authorization lasts, and whether the speaker receives additional compensation. They should

also say what happens to the model and source files when the agreement ends.

Privacy, biometric data, and AI rules

Voice recordings can reveal personal information. In some circumstances, voice data may be treated as biometric or sensitive personal data, particularly when it is processed to identify someone. Requirements can include clear notices, a lawful basis or consent, security measures, limits on retention, and a way to exercise privacy rights. The details depend on location and on what the system does with the audio.



AI laws are also changing. Some rules focus on disclosure when people interact with AI or encounter synthetic media; others address impersonation, deceptive content, or particular high-risk uses. Advertising and consumer protection rules matter too: a brand should not imply that a real person personally endorsed a product if they did not. Ask counsel to check current requirements for every market in which the mascot will be available.

What strong voice consent should cover

“We have permission” is not enough detail for a durable voice program. Good **voice consent** describes the actual use and makes it possible to audit what was approved. The speaker should understand that a model may generate lines they did not individually perform.

At a minimum, put these points in writing:

- **Purpose:** identify the mascot, product, and intended use—such as marketing, in-app narration, or customer support.
- **Scope:** list the channels, regions, languages, and types of content allowed. State whether the voice can answer open-ended questions or only play approved scripts.
- **Duration:** set a term and explain renewal, expiration, and what happens to existing content afterward.
- **Control:** describe review and approval rights, prohibited topics, and who can create or publish new lines.
- **Payment:** explain fees, residuals, usage-based compensation, or other agreed consideration.
- **Model handling:** specify who can access the voice model, whether a vendor processes it, and whether data may be used to improve other systems.

- **Retention and deletion:** define how long recordings, embeddings, and model files are kept, and what deletion means in practice.
- **Withdrawal and disputes:** explain how consent can be withdrawn, what uses can be stopped, and how existing ads, downloads, or cached files are handled.
- **Disclosure:** decide when the audience will be told that the voice is synthetic or AI-generated.

Consent should be informed and freely given—not buried in a broad release or inferred from a recording session. If a voice will be used in a new way, get fresh approval rather than stretching old wording beyond what the speaker reasonably expected.

What changed in AI voice quality and conversation handling?

Earlier phone systems often made users navigate rigid menus: “Press 1 for billing, press 2 for support.” Scripted chatbots can feel similar when they recognize only a narrow set of phrases and send people in circles after a slight change in wording. In support, a system that cannot handle interruptions, corrections, or a request to speak to a person quickly becomes another obstacle.

Modern speech models can sound more natural, and conversational systems can keep track of context across multiple turns. But smoother speech is not the same as trustworthy behavior. A polished voice can still misunderstand a request, invent an answer, or fail to hand off to a human. Voice products should be tested for interruptions, accents, background noise, corrections, silence, and escalation—not just for a pleasant demo.

One useful test is to speak over the agent mid-sentence. Can the system stop, listen, and respond to the new request? Can it recognize “No, I meant the other order” without restarting the entire flow? If not, the experience may be a prettier IVR rather than a genuinely useful voice agent. And if the voice is cloned, the system needs guardrails so it does not generate damaging, misleading, or off-brand statements in that recognizable voice.

How TTS, ASR, and language models work together

A voice agent commonly combines three technologies:

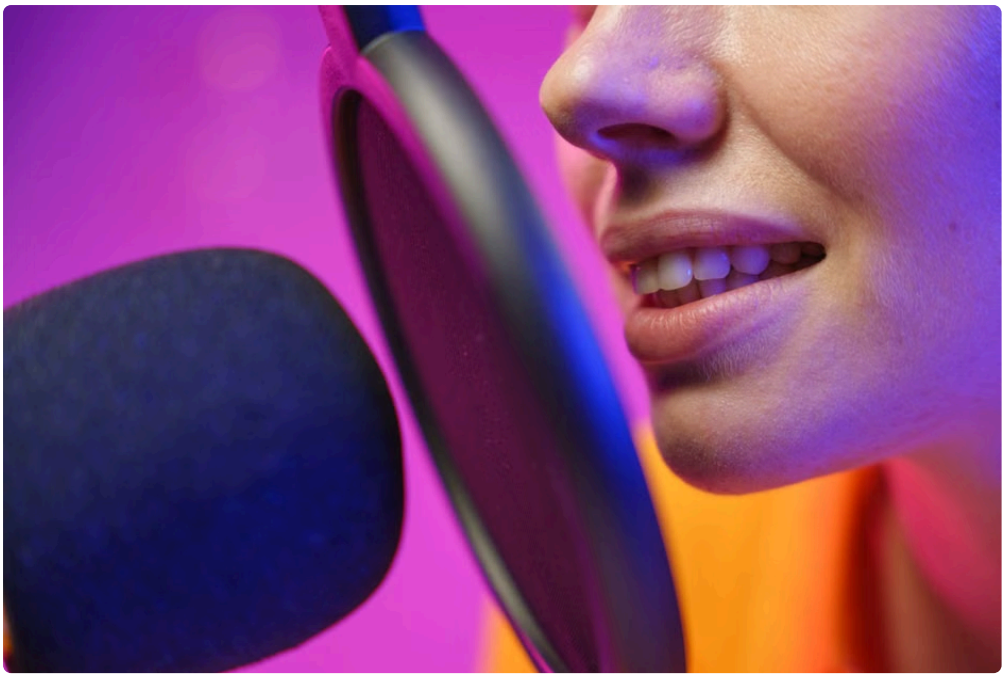
- **ASR (automatic speech recognition)** turns a person’s spoken words into text or another machine-readable representation.
- **A language model** interprets the request, uses conversation context, and selects or generates a response. In a support workflow, it may also call approved tools or retrieve information.
- **TTS (text-to-speech)** turns the response into spoken audio. A voice clone may be used at this stage to make the generated speech resemble an authorized speaker.

In practice, the components may overlap or operate in a more integrated real-time pipeline. The important point for rights management is that authorization for one part does not automatically authorize all the others. Permission to process speech for recognition is not necessarily permission to train a clone; permission to generate speech is not permission to use it in Extra resources every channel.

Each step also has failure modes. ASR can mishear names, numbers, or accents. A language model can produce inaccurate or inappropriate content. TTS can mispronounce a brand name or make an emotionally sensitive message sound wrong. The system needs clear boundaries, monitoring, and a reliable human handoff. Don’t promise a “human-like” experience without explaining what it can do, what it cannot do, and when a person takes over.

Why mobile apps are a natural channel—and a sensitive one

Mobile apps can provide useful context: a signed-in customer, a current order, saved preferences, or the screen they are already viewing. That context can reduce repetitive questions. The device also gives users a familiar place to start a conversation, whether by typing or speaking.



But identity and context make consent and privacy more important, not less. A voice interaction could involve account details or sensitive support topics. Tell users when they are speaking to an AI system, request microphone access only when needed, explain how audio is processed, and avoid collecting or retaining more than the service requires. Make it easy to mute, switch to text, or reach a human.

For example, the ElevenLabs iOS App may help creators explore voice-generation workflows on iOS. That is different from receiving legal permission to clone a real person for a commercial mascot. A tool's capabilities, settings, or terms do not replace the speaker's consent, the brand's rights review, or the need to meet applicable privacy and disclosure obligations. Check the provider's current terms and data-handling documentation before using any service with real people's recordings.

A practical launch checklist

1. **Identify the speaker and rights holders.** Confirm who can authorize the voice and whether a representative, estate, employer, or union has relevant rights.
2. **Define the use before recording.** Decide whether the mascot will read fixed copy, answer questions, appear in ads, or interact with customers.
3. **Obtain specific written consent.** Cover model creation, permitted outputs, channels, geography, duration, compensation, and withdrawal.
4. **Set rights management controls.** Limit who can generate audio, keep an approval trail, version the model, and block prohibited topics or claims.
5. **Review vendor and privacy terms.** Check subprocessors, training use, retention, deletion, security, and whether the provider can retain or reuse recordings.
6. **Test the whole conversation.** Try interruptions, corrections, noisy environments, out-of-scope requests, and human handoffs. Include tests for misleading endorsements and unsafe outputs.

7. **Plan disclosure and support.** Tell users when the mascot is synthetic where appropriate, and provide a clear route to a human.
8. **Revisit the approval.** A new language, campaign, channel, or interactive feature may fall outside the original consent.

Question Why it matters What to document Who approved the clone? The speaker may not be the only relevant rights holder. Identity, authority, releases, and any union or representative approvals. What can it say and where? Broad, open-ended use creates greater legal and brand risk. Purpose, channels, topics, locations, duration, and approval process. What happens to the voice data? Recordings and models raise privacy, security, and retention questions. Access, vendor processing, training use, retention, deletion, and incident response. How will users understand the interaction? A synthetic voice can be mistaken for a real person or endorsement. Disclosure language, escalation path, and rules for advertising claims.

The bottom line

You may be able to clone a real person's voice for a brand mascot legally, but the responsible answer starts with permission—not the software. Get explicit, specific voice consent; confirm the relevant publicity, contract, privacy, and AI requirements; and maintain practical rights management over who can generate what. Then test the experience as rigorously as you would any customer-support flow.

If the brand cannot explain what the speaker agreed to, how long the model will exist, how users will know it is synthetic, and how a person can take over when the agent fails, it is not ready to launch.