“We support SIP” can mean anything from “you can connect a carrier with a supported setup” to “we have a SIP option somewhere in the enterprise plan.” Before choosing a voice agent platform, verify what SIP trunking support actually includes: who supplies the phone numbers, how calls reach the agent, how failures are handled, and whether you can control routing and call behavior.

The quickest way to find out is to test the entire phone deployment—not just the agent’s conversation in a polished demo. Call it from a quiet room, then call again with background noise. Follow a realistic call through your phone lines, the telephony integration, the voice agent, and any handoff to your support system. Listen for delays, interruptions, and missing context.

This guide provides a practical checklist for evaluating SIP trunking support, call routing, latency, timeouts, and knowledge quality. It applies whether you are connecting an existing carrier, adding voice to a support workflow in Zendesk, or comparing a phone-based deployment with a website widget.



What SIP Trunking Support Should Mean

SIP trunking is a way to carry phone calls over an IP network between your telephony provider and another system. In a voice agent setup, a SIP trunk may connect your carrier’s infrastructure to a voice platform, a communications provider, or a call-routing layer that sends calls to the agent.

But a platform’s SIP trunking support can refer to several different arrangements:

- **Bring your own carrier:** You connect a carrier or PBX that you already use.
- **Platform-provided telephony:** The vendor supplies numbers and handles the phone connection.
- **Partner integration:** SIP connectivity is available through a named telephony partner, with its own setup and support boundaries.
- **Custom or enterprise setup:** SIP may be possible, but only with a particular plan, contract, or implementation project.

These options are not interchangeable. Ask the vendor to draw the call path and identify who owns each part: the carrier, SIP trunk, voice platform, agent, and destination system. If the answer is vague, you do not

yet know what “supported” means.

Map the Voice Agent Pipeline Before You Buy

A live call passes through several stages. Any one of them can introduce delay or break the experience, so evaluate the complete pipeline rather than only the model or voice.

1. **Inbound call:** A caller dials a number on your phone lines. The carrier receives the call and applies its initial routing rules.
2. **SIP connection:** The carrier or PBX sends the call to the voice platform, directly or through a telephony provider.
3. **Call routing:** Rules determine whether the call goes to an agent, a queue, a person, a department, or another destination.
4. **Audio processing:** The system detects speech, converts it to text or otherwise interprets it, and produces a response.
5. **Knowledge retrieval:** If the answer depends on company information, the agent may retrieve it from a knowledge base or a retrieval-augmented generation (RAG) system.
6. **Speech response:** The agent’s answer is generated as audio and returned to the caller.
7. **Handoff and logging:** The system may transfer the call, create a ticket, or pass a summary and other context to Zendesk.

For each stage, ask what can fail and how the failure appears to the caller. A SIP connection could be rejected. A routing rule could send the call to the wrong queue. Speech recognition could struggle with noise. A knowledge lookup could return an outdated policy. A transfer could connect but lose the caller’s reason for calling.

Ask for the Call Path in Concrete Terms

Request a diagram or written description of a typical inbound call. Confirm whether the platform accepts SIP directly, requires a provider in between, or only supports a particular carrier. Ask which side configures the SIP trunk and where call routing rules live.

Also clarify the practical setup requirements. Depending on the architecture, these may include SIP addresses, authentication, network or firewall requirements, supported codecs, caller ID behavior, and rules for transferring calls. Do not assume that a generic compatibility statement guarantees your carrier, existing PBX, or routing design will work.

Use This Checklist to Verify SIP Trunking Support

Use the questions below in a vendor call. Ask for demonstrations or documentation, not just verbal assurances.



Area What to ask What to verify Connection Can we connect our own SIP trunk or carrier? Supported setup, dependencies, and who configures it Phone numbers Can we use our existing numbers? Number ownership, porting options, and caller ID behavior Call routing Where are routing rules managed? Rules for agent, queue, person, overflow, and after-hours calls Transfers Can the agent transfer a caller to a live person? Whether the destination receives relevant call context Failure handling What happens if the agent or connection is unavailable? Fallback destination, alerts, and recovery behavior Controls Can we configure timeouts, silence handling, and interruption? Visible settings and the ability to test them Support Who owns troubleshooting across the carrier and platform? Clear escalation path and useful call diagnostics

When you compare ElevenLabs, ElevenAgents, or another platform, compare the actual deployment model and controls—not just the names of supported integrations. Confirm which features are included in the plan you would buy, and whether your desired phone deployment has been tested with your carrier and call flows.

Test Call Routing and Failure Points

Call routing is where a technically working connection can still become a poor customer experience. Test ordinary calls, but also test the paths that are easy to overlook: after-hours calls, busy queues, unavailable agents, transfer failures, and callers who need a human immediately.

For each path, confirm:

- Which number or route receives the call first.
- What conditions send the call to the agent versus a human team.
- What happens if the SIP connection or agent is unavailable.
- Whether callers hear silence, a prompt, a queue, or a fallback destination during a failure.
- Whether a transfer preserves caller ID and passes a useful summary or other context.
- Where recordings, transcripts, and routing events are available for review.

A good escalation is not merely “transfer completed.” The receiving person should know why the caller is there and what the agent has already tried. If your workflow uses Zendesk, check whether the voice

interaction can be associated with the right customer or ticket and whether the handoff details are visible to the support team.

Evaluate Latency and Turn-Taking in Real Calls

Latency is the time between a caller finishing a thought and hearing a useful response. That interval can include speech detection, recognition, any RAG search, response generation, and audio playback. A vendor may report a low average while real callers experience awkward pauses, so test the deployed call path yourself.

Listen for whether the agent waits too long after the caller stops speaking, starts speaking before the caller is done, or takes a noticeable pause after a short question. Also test longer statements, corrections, interruptions, and brief silence while the caller thinks.

In my call reviews, I count dead-air seconds out loud because “fast” is not a useful measurement unless the conditions are clear. Record the call, note the time between turns, and test the same scenario more than once. Run one test in a quiet room and another with background noise. A studio-quality demo tells you little about a busy customer calling from a car, shop, or shared office.

Test Interruption Handling and Timeouts

Interruption handling, sometimes called barge-in, determines whether the agent can stop talking and listen when the caller speaks over it. Ask if this behavior can be configured and tested. A system that cannot be interrupted can trap callers in long responses; a system that interrupts too eagerly may cut off useful explanations.

Test awkward but realistic scripts, such as:

- “Actually, that’s not what I meant.”
- “Wait—can you transfer me to a person?”
- “I have the order number, but let me find it.”
- A caller beginning a correction while the agent is midway through a response.

Ask which timeout settings are available. These may cover how long the system waits for the caller to start speaking, how much silence it tolerates mid-sentence, and when it reprompts or ends a call. Make sure the settings are visible to your team rather than hidden behind a vendor-only process. Test what the caller hears when a timeout occurs; a silent cutoff is not a graceful recovery.

Check Knowledge Bases and RAG Quality

SIP trunking gets a call to the agent, but it does not make the agent’s answers correct. If the system uses a knowledge base or RAG, test whether it can find the right source and respond clearly over a phone connection.

Prepare questions with different levels of difficulty: a common policy question, a product detail, an exception, and a question whose answer is not in the knowledge base. Include information that is similar across multiple documents. Then check whether the agent retrieves the current ai call center agent review source, gives a consistent answer, and says when it does not know.

Evaluate the audio response as well as the underlying answer. A correct answer can still be hard to use if it is too long, full of web-style formatting, or delivered as a dense list over the phone. Ask how knowledge

sources are updated, whether indexing status is visible, and how an admin can investigate a poor answer. Include the same test questions in your website widget if you offer one, but do not assume that widget performance proves the phone experience works: voice adds turn-taking, noise, and time pressure.

Run a Practical Pilot Before Committing

Use a small, controlled pilot with real call flows. Include people who know the service and people who do not; first-time users often expose unclear prompts that internal testers overlook.

1. **Write test scenarios:** Include normal requests, interruptions, background noise, unsupported questions, and transfer requests.
2. **Test both ends of the connection:** Call from the quiet-room setup, then repeat with realistic noise and more than one phone if possible.
3. **Measure the experience:** Note answer delays, repeated questions, misheard details, timeouts, and failed routes.
4. **Review the handoff:** Confirm the human destination receives context and that the event appears in your support workflow.
5. **Repeat after changes:** Retest after editing prompts, routing, knowledge sources, or telephony settings.

Do not judge the pilot only by whether a call connected. Track whether callers reached the right destination, whether the agent recovered from interruptions, and whether the final answer or escalation helped resolve the reason for the call. Keep test conditions consistent so you can compare vendors or configuration changes fairly.

Red Flags to Watch For

- The vendor says “SIP supported” but cannot explain the call path or identify dependencies.
- There is no clear fallback when the agent, carrier connection, or destination is unavailable.
- Interruption and timeout behavior cannot be demonstrated or configured.
- Transfers do not pass the caller’s intent or previous conversation to the receiving team.
- The demo uses pristine audio only, with no realistic noise or awkward caller scenarios.
- Knowledge-base answers cannot be traced, tested, or corrected by your team.
- Responsibility for troubleshooting is unclear when an issue crosses the carrier and platform boundary.

Make the Decision on the Whole Experience

Checking SIP trunking support means verifying more than a box on an integration list. Confirm how your carrier connects, who owns call routing, which phone lines and numbers you can use, and what happens when a part of the pipeline fails. Then test latency, turn-taking, interruptions, timeouts, and knowledge quality in actual calls.

A voice agent platform is ready for phone deployment when the connection is clear, the routing is controllable, failures have sensible fallbacks, and callers can interrupt or reach a person without losing context. Whether you choose ElevenLabs, ElevenAgents, or another provider, insist on a test you can reproduce—not a promise that the system is “human-like.”