BUYER'S GUIDE · CONVERSATION INTELLIGENCE

15 questions to ask every conversation intelligence vendor.

Evaluating vendors for your patient services, hub operations, or patient access teams? The vendor landscape is crowded. Not every solution is what it appears to be in a demo. Use these questions to go beneath the surface and understand what you're actually buying.

WHY THIS GUIDE EXISTS

The vendor landscape is crowded. Not every solution is what it appears to be in a demo.

Conversation intelligence has become the must-have layer for healthcare contact centers, and the category has filled with vendors making nearly identical claims. Some are mature, healthcare-native platforms. Others are white-labeled resellers, generic tools retrofitted for healthcare, or ambitious roadmaps wrapped in a polished sales motion. The questions in this guide are designed to surface the difference. Use them on every vendor you evaluate, including us.

HOW TO USE THIS GUIDE

01 Ask every question to every vendor, including any incumbents you are re-evaluating.

02 Use the follow-up questions to push past rehearsed answers and get to specifics.

03 Use the scoring sheet at the back to evaluate vendors consistently across your team.

PART 01

Technology & AI.

What you are actually buying and how the AI was built.

01

OF FIFTEEN

WHAT IS THIS SOLUTION BUILT TO DO?

"What is this solution's core business focus, and what is it actually built to do?"

The most important question is whether the vendor is selling you both a solution and the means to monitor whether it's working.

Understand what the vendor's primary business is and what drove the decision to build this solution. A platform built from the ground up for conversation intelligence is a fundamentally different product than one where it was added as a feature. Ask how accuracy is established on day one, how it improves over time, and how it is maintained, and whether you can benchmark results against your own data to verify claims independently.

FOLLOW-UP QUESTIONS

- → Is conversation intelligence your core product, or a capability added to an existing platform?
- → **What is your day-one accuracy, how does it improve over time, and how is it maintained?**
- → Can we benchmark your accuracy results against our own data to independently verify performance claims?
- → Are you also providing the monitoring framework used to evaluate whether your solution is working?

02

OF FIFTEEN

AI MODELS & TRAINING

"Is this solution classified as AI or AI-assisted — and how are the underlying models built, trained, and monitored?"

The architecture of the AI matters as much as the outputs it produces.

Not all "AI" in this category means the same thing. Some vendors apply general-purpose language models to conversation data. Others have purpose-built, proprietary models trained on healthcare-specific information. Ask whether the AI is built in-house or sourced from a third party, how training data is curated and validated, and how model outputs are monitored for accuracy over time.

FOLLOW-UP QUESTIONS

- → Is this solution classified as AI or AI-assisted, and what is the practical difference in your platform?
- → Are the models proprietary and built in-house, or sourced from a third-party provider?
- → **What data is used to train and calibrate the models, and how are results monitored for quality and accuracy?**
- → Does the solution include an in-app AI assistant, and what documentation or data does it rely on?
- → How do you handle model drift, and how frequently are models retrained?

03

OF FIFTEEN

DATA INGESTION & TRANSCRIPTION

"Are you pulling in raw audio or transcriptions directly from the CCaaS vendor?"

If the transcription is wrong, every downstream output is affected.

Healthcare conversations are complex and full of nuance. Clinical terminology, benefit language, coverage questions, and member or patient context all require specific knowledge that many generic transcription tools miss. Every AI model downstream runs on top of that transcription layer, so inaccuracy compounds quickly. Vendors with their own transcription engine built for healthcare have a meaningful accuracy advantage. Ask for benchmark data specific to your use case.

FOLLOW-UP QUESTIONS

- → **What is your transcription accuracy rate, and how do you measure it?**
- → Have you tested accuracy specifically against healthcare terminology?
- → Can you re-transcribe existing audio, or are you locked into what the CCaaS provides?

04

OF FIFTEEN

CCAAS INTEGRATIONS

"What out-of-the-box integrations do you have with leading CCaaS platforms?"

Getting data in needs to be easy, accurate, and reliable.

What matters most is that your conversations flow into the platform reliably, completely, and with minimal friction on your team. Whether that is a direct native connection, a file-based workflow, or another method, the key questions are around completeness, consistency, and what happens when something breaks. Ask specifically about the platforms your organization runs and how the vendor handles gaps or errors in data delivery.

FOLLOW-UP QUESTIONS

- → **Is your integration with our CCaaS platform native, or does it require middleware?**
- → How do you handle multi-channel ingestion: voice, chat, and email?
- → What does a standard implementation timeline look like for a CCaaS connection?

05

OF FIFTEEN

HEALTHCARE & INDUSTRY SPECIFICITY

"Was this purpose-built for healthcare, or adapted from a general-purpose platform?"

Generic tools retrofitted for healthcare consistently lag native ones.

General-purpose conversation intelligence, built for sales or broad contact center use, has to be retrofitted for healthcare. The language, compliance requirements, and use cases are fundamentally different. Adapted solutions trail native ones in model accuracy, taxonomy depth, and regulatory readiness. Ask specifically about healthcare: what share of their customer base is in healthcare, and whether their roadmap is shaped by the specific requirements of healthcare organizations.

FOLLOW-UP QUESTIONS

- → **What percentage of your customer base is in healthcare?**
- → How does your taxonomy reflect workflows specific to our segment, prior auth and hub intake for pharma, or claims, appeals, and member services for payers?
- → Do you have proprietary models built for healthcare compliance, including automated HIPAA identification, adverse event detection, or payer-specific regulatory requirements?
- → What percentage of your roadmap is focused on healthcare-specific requirements versus general platform development?

PART 02

Product capabilities.

What the platform actually does for your team.

06

OF FIFTEEN

CUSTOMER CENTRICITY

"Does this solution focus exclusively on agent performance, or does it surface friction in the broader customer journey?"

Agent performance and customer experience are not the same problem.

Many conversation intelligence tools were built to evaluate agents: scoring calls, flagging compliance issues, coaching reps. The more valuable question is whether the platform also surfaces what is driving customer friction and dissatisfaction. For pharma organizations that may mean therapy abandonment or access barriers; for payers and providers it may mean member dissatisfaction, call avoidance, or unresolved service failures. Ask whether the solution can ingest 100% of conversations, analyze unsupervised topics, and proactively surface trends increasing call volume or service costs.

FOLLOW-UP QUESTIONS

- → Does the solution surface friction in the customer journey, or is it focused on agent performance only?
- → **Can the solution ingest 100% of all sources of conversations, including from other teams, hubs, chat bots, call centers, and more, or does it rely on sampling?**
- → Does the solution identify issues specific to your business, therapy adherence and access for pharma, or member satisfaction and service gaps for payers and providers?
- → Can the platform uncover unknown or emerging issues, topics you have not tagged or labeled in advance?
- → Does the solution proactively identify trends driving call volume or service costs?

07

OF FIFTEEN

AGENT COACHING & QUALITY ASSURANCE

"What agent coaching and QA capabilities does this solution support, and how much is automated?"

The best platforms support agents and monitor quality across 100% of conversations, not just the ones a reviewer happens to pull.

A capable platform should evaluate agent performance and call quality automatically, not just flag recordings for a human reviewer. Equally important is whether the platform uses objective, consistent rubrics so agents understand exactly what they are being evaluated against and what good looks like. Inconsistent or opaque scoring criteria undermine agent trust and make coaching less effective. Ask whether the solution supports AI-automated QA scoring and structured reporting by team, call line, and call center.

FOLLOW-UP QUESTIONS

- → Can the solution evaluate agent performance and call quality without requiring manual review?
- → **Is QA scoring fully automated, partially automated, or primarily human-driven?**
- → Can you report on performance at the team, call line, and call center level?
- → How is coaching delivered to agents, and what does the agent-facing experience look like?

08

OF FIFTEEN

CUSTOMIZATION & CONFIGURABILITY

"What control do you have over models, evaluations, and data structure, and how much requires vendor involvement?"

Flexibility in the platform determines how fast you can respond to new business questions.

Some platforms lock you into a fixed taxonomy and evaluation structure, requiring a services engagement every time your needs change. Others give you direct control: custom evaluations, no-code model deployment, keyword tags, and configurable data hierarchies. The degree of configurability you have determines how quickly your team can pivot when new compliance requirements emerge or call patterns shift.

FOLLOW-UP QUESTIONS

- → Can you bring your own custom evaluations to the platform without vendor build time?
- → **Can you deploy custom, no-code models within the platform, not reliant on keyword libraries?**
- → Do you have the ability to set up custom keyword tags that are reportable?
- → What control do you have over the structure of your data hierarchies within the platform?
- → How are configuration changes managed: self-service, or through a vendor services request?

09

OF FIFTEEN

SPEED TO VALUE

"What out-of-the-box models do you have, and how long until you deliver insights?"

Some vendors are selling a vision. The price tag is a multi-quarter services engagement.

Many newer vendors require 6–12 months of custom model build before any insights surface, which is less a software purchase and more a professional services engagement. A strong out-of-the-box model library compresses time-to-value dramatically. Understand clearly what you can access on day one versus what requires a lengthy configuration period.

FOLLOW-UP QUESTIONS

- → **What can we activate on day one, versus what requires custom configuration?**
- → What is the average time to implement and begin seeing results?
- → Can you walk us through an actual client go-live timeline, start to finish?

PART 03

Risk & compliance.

Security, business viability, and proven results.

10

OF FIFTEEN

SECURITY & PRIVACY

"How do you handle PHI, HIPAA, and data residency — and have you completed a third-party security audit?"

Compliance is a spectrum, not a checkbox.

Patient services calls contain protected health information by definition. Any vendor here must be HIPAA-compliant, but compliance varies widely. Ask for their BAA, data residency policies, audit history, and breach notification process. Beyond HIPAA, ask about GDPR and CCPA compliance. Strong vendors support automated redaction of personal identifiers and PHI, offer bring-your-own-encryption-key options, and use SCIM for secure account management.

FOLLOW-UP QUESTIONS

- → Are you prepared to sign a BAA, and what does your standard BAA cover?
- → **When was your most recent SOC 2 Type II audit, and can you share the report?**
- → Is the solution GDPR and CCPA compliant in addition to HIPAA?
- → Can the solution automate redaction of personal identifiers and protected health information?
- → Do you offer bring-your-own-encryption-key (BYOK) and SCIM user management?

ROADMAP VIABILITY & BUSINESS RISK

"What happens to this product if adoption is slower than projected?"

The risk isn't that the product doesn't improve. It's that it gets discontinued.

New solutions carry real risk. If a vendor hasn't achieved sufficient adoption, they face a hard decision: continue investing or sunset the product. Ask about their investment thesis, customer acquisition targets, and 12-, 24-, and 36-month roadmap. If they cannot answer clearly, the roadmap may be aspirational rather than funded.

FOLLOW-UP QUESTIONS

- How long has this specific solution been in market, and how many paying customers do you have today?
- What are your customer acquisition targets for the next 12 months, and how are you tracking?
- If this product line underperforms, what is the organization's commitment to continuing investment?
- **Can you share your 18-month roadmap, and which items are funded vs. dependent on future revenue?**

12

OF FIFTEEN

PROVEN DEPLOYMENT

"Do you have live customers using this solution today, and can we speak with them?"

Pilots and POCs do not tell the whole story. Only live deployment at scale does.

The market is full of vendors, especially those entering from adjacent spaces like hub operations software, showing polished demos of capabilities not yet in production. The question is whether this has been deployed, used by real agents, and validated with paying customers in your industry. If they cannot offer a reference call or a named case study, that tells you something.

FOLLOW-UP QUESTIONS

- → Can you share a case study with measurable outcomes from a live deployment?
- → Do any of your current customers operate in our specific segment, pharma, payer, or provider?
- → **Would any of your customers be willing to speak with our team directly?**

PART 04

The partnership.

Who is actually accountable for your success after the contract is signed.

13

OF FIFTEEN

ROADMAP TRANSPARENCY

"What is generally available today versus on the roadmap — and what did we just see in the demo?"

Demos are not products. Future-state capabilities are routinely shown as if they exist today.

Understanding the line between what is live and what is in development protects you from buying a vision and waiting 12–18 months for the product to catch up. Ask explicitly. Get it in writing if a key capability is driving your decision.

FOLLOW-UP QUESTIONS

- → **What features in today's demo are live and available for all customers right now?**
- → What is in beta or limited release, and what is the GA timeline?
- → How do customers currently influence your roadmap prioritization?
- → What percentage of your roadmap is focused on customers' custom requirements vs. general platform development?

14

OF FIFTEEN

DEDICATED TEAM & INTERNAL PRIORITIZATION

"Is there a team solely focused on this solution, or is it shared with your existing product line?"

If their core product has a critical issue, who pauses your work?

This exposes a common risk with hub vendors who added conversation intelligence as a new capability. If the same engineers, PMs, and support staff manage both the existing platform and this new solution, you compete internally for attention. A dedicated team signals genuine organizational commitment. A shared team signals a side project.

FOLLOW-UP QUESTIONS

- How many engineers, PMs, and data scientists are dedicated exclusively to this solution?
- **Is the roadmap for this product managed independently, or does it compete with your core product's roadmap?**
- How is this team structured, and who owns the P&L for this specific product?

15

OF FIFTEEN

IMPLEMENTATION & ONGOING SUPPORT

"What do implementation and ongoing support actually look like, and who is doing it?"

There is a significant difference between a vendor that walks alongside you and one that hands you a login.

Self-service or train-the-trainer models can work for simple tools, conversation intelligence is not a simple tool. Model configuration, taxonomy alignment, integration troubleshooting, and ongoing insight delivery all require expertise. Ask specifically who is responsible for your success post-contract: a named CSM, a shared queue, or a community forum.

FOLLOW-UP QUESTIONS

- What does the implementation process look like, and who leads it on your side?
- Is implementation self-service, train-the-trainer, or do you provide a dedicated implementation resource?
- **How long after implementation can we see insights on conversations?**
- Post go-live, who is our primary point of contact, and how many accounts do they manage?
- How do you handle escalations, and what are your SLA commitments for issue resolution?

WHAT STRONG ANSWERS LOOK LIKE

A vendor worth partnering with answers these questions directly. With specifics. Not deflection.

They should be able to name their underlying technology stack, show you transcription accuracy benchmarks, connect you to a live reference customer, and clearly distinguish what is in production today from what is on the roadmap. If answers are vague, escalated to legal, or deferred to a follow-up, take note. That pattern in the sales process typically reflects the product itself.

SIGNS OF A STRONG ANSWER

- ✓ Specific data and named benchmarks, not general claims
- ✓ Reference customers offered proactively, not reluctantly
- ✓ Clear distinction between live capabilities and roadmap items
- ✓ Named implementation lead with a defined account capacity

SIGNS TO TAKE NOTE OF

- ! Vague answers deferred to a follow-up email
- ! No named reference customers in your industry
- ! Security questions escalated to legal without specifics
- ! Cannot clearly separate what is live from what is coming

VENDOR SCORING SHEET

Score each vendor against all 15 questions.

VENDOR _____ DATE _____

SCORE KEY: (S) Strong (M) Mixed (C) Concern (?) Not answered

PART 01 · TECHNOLOGY & AI

#	QUESTION	SCORE
01	What is the solution's core focus and day-one accuracy?	○
02	How are AI models built, trained, and monitored?	○
03	Raw audio or CCaaS transcriptions?	○
04	Native CCaaS integrations?	○
05	Purpose-built for healthcare and pharma?	○

SECTION NOTES

PART 02 · PRODUCT CAPABILITIES

06	Does it surface customer friction, not just agent performance?	○
07	What agent coaching and QA capabilities are automated?	○
08	What customization does the platform support?	○
09	Out-of-the-box models available on day one?	○

SECTION NOTES

PART 03 · RISK & COMPLIANCE

#	QUESTION	SCORE
10	HIPAA, GDPR, CCPA, SOC 2 Type II, redaction, BYOK, SCIM?	○
11	What happens if adoption is slower than projected?	○
12	Live paying customers in healthcare we can speak with?	○

SECTION NOTES

PART 04 · THE PARTNERSHIP

13	What's live vs. what's on the roadmap?	○
14	Dedicated team or shared with core product?	○
15	Who leads implementation and owns post-go-live support?	○

SECTION NOTES

STRONG	MIXED	CONCERN	OVERALL RECOMMENDATION
/ 15	/ 15	/ 15	