




AUTO ATTENDANT / RESEARCH & PRACTICAL GUIDANCE

AI Receptionist vs Answering Service vs Auto Attendant

Get a business phone number with a simple auto attendant menu from Auto Attendant. A caller picks who they need and the call rings that person's own mobile, so nobody on your team has to publish a personal number.

Published by Auto Attendant · 2026-09-22 · ai receptionist · answering service · auto attendant · small business phone system · virtual receptionist · call routing · contact centre · phone answering solution

Original article: <https://autoattendant.io/articles/ai-receptionist-answering-service-auto-attendant>



In this article

1. [Executive Summary](#)
 2. [Introduction and Background](#)
 3. [Auto Attendant](#)
 4. [AI Receptionist](#)
 5. [Human Answering Service](#)
 6. [Full Phone and Contact-Centre Platforms](#)
 7. [Feature Comparison](#)
 8. [Performance and Benchmarks](#)
 9. [Data Analysis and Evidence](#)
 10. [Implications and Future Directions](#)
 11. [Frequently Asked Questions \(FAQs\)](#)
 12. [Conclusion](#)
-

Executive Summary

The choice among an **AI receptionist**, **human answering service**, and **auto attendant** begins with the job after “hello.” An auto attendant asks the caller to choose a destination and routes the call. An AI receptionist conducts a bounded conversation, which can include answering configured questions, collecting details, scheduling, and transferring. A human answering service assigns an outsourced person to speak for the business. A full phone or contact-centre platform adds queues, agent controls, recording, reporting, and often multiple channels. Microsoft’s documentation captures the first boundary: an auto attendant lets callers “navigate a menu system” to reach the right department (learn.microsoft.com). AI receptionists differ because they accept natural-language requests rather than only predefined choices.

For a 2 to 20-person business, the neutral default is straightforward. Choose an **auto attendant** when callers mainly need the correct person and that person is available to answer. Choose **AI** when repeatable conversations, intake, appointment actions, or after-hours coverage create value, but define a fast human escape route. Choose a **live answering service** when calls are emotional, ambiguous, exception-heavy, or costly to mishandle. Choose a **full platform** when the operating requirement is a queue, supervisor visibility, recordings, workforce controls, or cross-channel service. Amazon documents that a contact-centre queue can hold contacts when no agent is free (docs.aws.amazon.com), while Microsoft exposes waiting-call, average-wait, and longest-wait measures (learn.microsoft.com).

Price is a shape, not a single sticker. As retrieved on **September 22, 2026**, autoattendant.io lists **\$29 per month for the whole company**, without per-user seats (autoattendant.io). [Smith.ai](https://smith.ai) lists a human Starter plan at **\$300 for 30 calls**, with **\$11.50 per additional call** (smith.ai). [My AI Front Desk](https://myaifrontdesk.com) lists AI call answering at **\$99 monthly**, or **\$79 monthly billed annually with 200 voice minutes** (myaifrontdesk.com). These are examples of different billing units, not universal category averages.

The safest selection method is a **20-call local pilot**, not a vendor-wide benchmark. Include ordinary calls, accents and noise, interruptions, urgent requests, ambiguous questions, transfer failures, and accessibility cases. Score completion, correct escalation, summary accuracy, recovery, and staff cleanup time. NIST

recommends realistic test sets and human intervention when an AI cannot detect or correct errors ([airc.nist.gov](https://www.airc.nist.gov)) ([airc.nist.gov](https://www.airc.nist.gov)). For clinics, recordings and transcripts can create Health Insurance Portability and Accountability Act (HIPAA) obligations, including a business associate agreement where applicable ([hhs.gov](https://www.hhs.gov)).

Introduction and Background

“Phone answering system” is an imprecise buying label. It can mean a press-key menu, conversational software, an outsourced person, or a complete operational system for managing agents and queues. Treating those as interchangeable produces the most expensive failure in this decision: buying sophisticated conversation when a transfer was enough, or buying simple routing when the caller needed someone to resolve an exception.

This report evaluates categories, not a ranking of AI vendors. The target reader is an owner of a small trade, clinic, property firm, or professional-services business. In these settings, one call may be a routine request for hours, another a booking, another a distressed customer, and another an urgent after-hours escalation. The correct system depends less on how advanced the technology sounds than on the **cost of a bad handoff**.

The basic terms should remain separate:

- **Auto attendant:** A deterministic menu, usually keypad or speech choices, that sends the caller to a person, group, voicemail, or another menu. Microsoft’s configuration example uses “For Sales press 1. For Support press 2” (learn.microsoft.com).
- **AI receptionist:** Software that interprets natural language and can respond or act within configured instructions. Current first-party examples document frequently asked question answers, intake questions, appointments, transfers, transcripts, and summaries (ringcentral.com).
- **Human answering service:** People outside the buyer’s company who answer in the company’s name and follow instructions. Ruby documents transfers, frequently asked questions, and messages as included call jobs (ruby.com).
- **Full phone or contact-centre platform:** A system for operating communication at team scale, including queues, routing profiles, reporting, recording, supervision, and sometimes chat, email, text, or video. Amazon’s event model covers voice, chat, task, and email contacts (docs.aws.amazon.com).

The practical question is therefore: **What must reliably happen after the greeting, and who is available when the system transfers the call?**

Auto Attendant

Capabilities

An **auto attendant** is a routing layer. It gives callers a small set of explicit choices, then executes the selected route. Typical destinations include one employee, a ring group, voicemail, an on-call number, or a second menu. It can also apply business-hours logic. In Microsoft’s phone-system documentation, a call queue is a

separate capability that can hold calls and find an available agent ([learn.microsoft.com](#)) ([learn.microsoft.com](#)). That distinction matters: menu routing is not automatically queue management.

Route-only service fits when intent is easy to express as a short branch:

- **New work:** Ring the estimator or sales mobile.
- **Existing appointment:** Ring dispatch or the office.
- **Billing:** Ring the bookkeeper or send to voicemail.
- **After hours:** Play a different greeting, reach the on-call person, or collect voicemail.

This is the category in which [autoattendant.io](#) belongs. Its current first-party page says one menu can have **up to nine options** and an unlimited number of people receiving calls ([autoattendant.io](#)) and explicitly says it does not answer the customer's phone on the customer's behalf ([autoattendant.io](#)).

Adoption

For a small team, adoption can be operationally light because employees may continue receiving ordinary calls on existing mobiles. The owner still has design work: decide the greeting, limit the menu depth, assign each option, establish an unanswered-call path, and maintain the after-hours schedule. A menu with nine possible options does not mean nine should be presented at once. Accessibility guidance favors a visible route to a person: the World Wide Web Consortium (W3C) says a phone menu option should be reserved for reaching one ([w3.org](#)).

Strengths and Limitations

The strengths are predictability, low configuration complexity, and a clean handoff to the owner's own team. The limitation is equally clear: it cannot conduct a genuine intake conversation, interpret an unanticipated situation, reconcile a calendar, or calm a distressed caller. If no one accepts the transfer, routing has not solved answering. A dead-end transfer is therefore the defining failure mode.

Choose this category when the work is **identify intent, then route**. Do not choose it merely because the product page uses "virtual receptionist" as a synonym. Category should follow behavior, not marketing language.

AI Receptionist

Capabilities

An **AI receptionist** adds conversation and action to routing. Current products can ask configurable intake questions, collect names and other details, consult approved content, book or reschedule appointments through a connected calendar, and pass a summary to the transfer recipient ([ringcentral.com](#)) ([ringcentral.com](#)) ([ringcentral.com](#)).

That capability is useful where calls are numerous but structurally repetitive:

- **Frequently asked questions:** Hours, location, service area, basic eligibility, or document requirements.

- **Lead intake:** Name, address, problem type, desired timing, and permission to text.
- **Scheduling:** Book, change, or cancel within explicit calendar rules.
- **Qualification:** Apply declared criteria, while avoiding discretionary decisions the system was not designed to make.
- **Transfer:** Route with context, rather than forcing the caller to repeat everything.

Adoption

Adoption is a knowledge-management project. The buyer must specify what the system may say, what it must never infer, how it authenticates a caller, which integrations can write data, and exactly when it escalates. The operating owner should be named, because policies, prices, coverage areas, and calendars change.

Recording and transcription require particular attention. RingCentral says its product can automatically record and transcribe every AI receptionist call (ringcentral.com). Its July 2026 privacy datasheet also says third-party AI services supply some features and that the greeting can be customized to notify callers they are interacting with AI (assets.ringcentral.com) (ringcentral.com). Those statements illustrate questions every buyer should ask, not conclusions about all AI services.

Strengths and Limitations

AI can provide immediate, consistent coverage and complete bounded tasks without consuming an employee's time. It is weakest when the caller's meaning depends on emotion, tacit context, unusual policy, conflicting instructions, or a situation absent from its approved material. Speech recognition performance also varies across speakers. A peer-reviewed 2020 benchmark reported aggregate word error rates of **0.35 for Black speakers and 0.19 for white speakers** across five commercial systems (pnas.org) (pnas.org). That study is not a current product score, but it is a reason to test with the actual caller population.

The required control is a short path to a person. Smith.ai itself notes that scenarios requiring a human receptionist arise (smith.ai). NIST likewise recommends human intervention when the system cannot detect or correct errors (airc.nist.gov).

Human Answering Service

Capabilities

A **human answering service** places a trained external receptionist between the caller and the business. The person follows a script, answers approved questions, takes a message, transfers the call, schedules, and may perform structured intake. Ruby documents integrations with a customer's existing calendar system (ruby.com) and says receptionists can enter customer details into a payment-processing link (ruby.com).

The differentiator is not that a human can do everything. It is that a person can recognize ambiguity, adjust tone, ask a clarifying question, and decide that the written script does not fit. Human service is therefore a strong candidate for bereavement, anxious patients, urgent property problems, angry clients, complicated scheduling, and other calls where the interaction itself is part of the service.

Adoption

Implementation still requires precise rules. The service needs an approved greeting, business knowledge, transfer destinations, schedules, emergency instructions, payment boundaries, and a feedback loop for incorrect handling. Smith.ai describes its human service as covering complex scheduling, detailed notes, and sensitive conversations (smith.ai). It also documents human escalation for advanced coordination or payment-sensitive information (smith.ai).

Strengths and Limitations

Human judgment and conversational repair are the principal strengths. Constraints include cost, capacity, service consistency, the precision of the client's instructions, and whether an actual live fallback is included at the hours needed. A buyer should verify if billing begins when the receptionist answers, whether transfers and brief spam calls count, and whether time is rounded.

AnswerConnect, for example, states that interactions are rounded up and billed in one-minute increments, while the first **30 interactions shorter than 30 seconds** in a billing cycle are not charged (answerconnect.com) (answerconnect.com). Those terms show why “per minute” must be defined before forecast costs are compared.

Full Phone and Contact-Centre Platforms

Capabilities

A **full phone platform** becomes the relevant category when the business needs to operate a queue, not merely answer a line. Amazon documents routing based on an agent's profile, queue hours, and configured flow logic (docs.aws.amazon.com) (docs.aws.amazon.com). It also documents skills-based routing and multiple queues (docs.aws.amazon.com) (docs.aws.amazon.com).

Adoption

This category requires more administration: user identities, numbers, devices or apps, permissions, routing profiles, retention settings, dashboards, integrations, and supervisor practices. Microsoft documents up to **45 days** of historical metrics for its queue tooling (learn.microsoft.com) and supervisor functions including listening and private coaching (learn.microsoft.com). Genesys documents workforce functions covering coaching, scheduling, performance, and quality management (genesys.com).

Strengths and Limitations

The platform is justified when managers need evidence of service levels, supervisors need control, calls must wait in an organized queue, or voice must join chat, text, email, or video. Zoom's contact-centre materials describe common routing for voice, video, chats, and messages, plus customer relationship management and workforce integrations (media.zoom.com).

Its weakness for a very small company is overhead. Per-user pricing, implementation effort, ongoing administration, and features that nobody uses may outweigh the operational benefit. If the real need is “press 1, then ring Dana's mobile,” a contact centre is the wrong category even if its feature list is impressive.

Feature Comparison

Table 1 maps call jobs to categories. “Verify” means the capability exists in some offerings but must be confirmed for the chosen plan, integration, hours, and jurisdiction.

Call job or requirement	Auto attendant	AI receptionist	Human answering service	Full phone/contact centre
Identify intent	Can do only through fixed menu choices.	Can interpret bounded natural language.	Can ask and clarify.	Verify, depends on attendant, interactive voice response, AI, or agent design.
Route to the owner’s team	Core job. autoattendant.io: \$29/month company-wide, up to nine options, unlimited receivers (autoattendant.io).	Common capability, verify context passed.	Common capability, verify transfer billing and availability.	Core capability with queues and routing profiles.
Answer routine questions	Cannot do beyond recorded information.	Can do from approved knowledge.	Can do from client instructions.	Verify the selected AI, interactive voice response, or agent module.
Qualify or collect intake	Limited to menu selection.	Can do structured intake, verify fields and write access.	Can do scripted or adaptive intake.	Can do through agents, forms, or automation.
Schedule	Cannot negotiate a calendar.	Verify calendar integration, permissions, and conflict rules.	Verify calendar access and script.	Verify product module and integration.
Take payment details	Cannot do.	Verify payment scope, authentication, security, and human fallback.	Verify script, system, and agent authorization.	Verify payment workflow and compliance boundary.
Handle distress or ambiguity	Cannot converse, provide direct human option.	Escalate early.	Best category candidate where judgment and tone dominate.	Agents can handle, provided staffing and routing are configured.
Queue callers	Usually not by itself.	Verify whether queueing is native or a transfer dependency.	Provider manages its own answering pool, but buyer-side queue behavior varies.	Core function: hold, priority, skills, overflow, and reporting.
Record, transcribe, analyze	Usually voicemail only, verify.	Often available, verify consent, retention, and subprocessors.	Verify recording, notes, retention, and access.	Common, with administrative and supervisor controls.

The matrix makes one pattern visible: “can answer” is not one requirement. A route-only call, a deterministic booking, a frightened patient, and a ten-person support queue are four different operating problems. The best phone answering solution for a small business is the smallest category that safely completes its real call jobs

and has a defined recovery path.

A practical decision tree

Apply these questions in order:

1. **Does the caller need a conversation?** If no, use an auto attendant and route to the team.
2. **Is the conversation repetitive, bounded, and testable?** If yes, pilot AI against the actual workflow.
3. **Are ambiguity, emotion, discretion, or exception handling central?** If yes, evaluate a human answering service.
4. **Must callers wait for the next available employee?** If yes, evaluate call queues or a full platform.
5. **Are recordings, live dashboards, supervision, workforce scheduling, or multiple channels required?** If yes, the requirement has moved beyond a standalone receptionist.
6. **Who receives the escalation?** If nobody is available, an advertised transfer feature is not a complete fallback.

Wrong-category signals are equally useful. Repeated dead-end transfers indicate a staffing or queue problem. Repeated incorrect policy answers indicate a knowledge-control or human-judgment problem. Employees manually retyping every summary indicate an integration problem. Managers asking how many callers abandoned the line indicate a reporting and queue problem.

Performance and Benchmarks

Public feature lists do not establish performance on a specific company's calls. No credible universal benchmark shows that one category always produces better outcomes. Performance should therefore be measured as **successful work**, not as whether a voice sounded natural.

The 20-call pilot

A 20-call exercise is a practical local screen, not a statistically validated sample size. Build it from real call types and run the same cases through finalists. NIST recommends testing before deployment and regularly during operation ([airc.nist.gov](https://www.nist.gov/airc)). It also advises varied scenarios, including crises and ethically sensitive contexts ([nvlpubs.nist.gov](https://www.nist.gov/nvlpubs)).

Use this call mix:

- **Five ordinary calls:** Common questions, standard bookings, and direct transfers.
- **Four exception calls:** Unsupported request, schedule conflict, policy edge case, and unavailable recipient.
- **Three noisy calls:** Road noise, speakerphone, and weak connection.
- **Three varied-speaker calls:** Different accents, speech rates, or vocabulary used by actual customers.
- **Three high-risk calls:** Distress, urgent after-hours need, and payment-sensitive request.

- **Two accessibility calls:** Repetition request, speech disability, or alternative-channel need.

W3C says persistent voice-menu errors should lead to a human operator and that voice-dependent services need alternatives such as text chat for people with speech disabilities ([w3.org](https://www.w3.org/)) ([w3.org](https://www.w3.org/)).

Table 2 is a scorecard template. Leave pass thresholds blank until the owner defines the cost of each failure.

Calls	Expected action	Observed action	Escalation result	Summary accuracy	Recovery	Staff cleanup minutes
1 to 5, ordinary	_____	_____	_____	_____	_____	_____
6 to 9, exceptions	_____	_____	_____	_____	_____	_____
10 to 12, noise	_____	_____	_____	_____	_____	_____
13 to 15, varied speakers	_____	_____	_____	_____	_____	_____
16 to 18, high risk	_____	_____	_____	_____	_____	_____
19 to 20, accessibility	_____	_____	_____	_____	_____	_____

The most important fields are completion and recovery. A system that misunderstands but quickly reaches the right person may be safer than one that sounds polished and confidently completes the wrong action. NIST's older spoken-dialogue research found a relationship between word error rate and task completion, while also emphasizing strategies that accomplish tasks despite recognition errors ([nist.gov](https://www.nist.gov/)) ([nist.gov](https://www.nist.gov/)).

Data Analysis and Evidence

Compare total cost, not headline price

As of **September 22, 2026**, representative official pages demonstrate several incompatible pricing units. Smith.ai prices live reception by calls, while AnswerConnect describes minute rounding. My AI Front Desk combines a subscription, included voice minutes, and credits. Goodcall says it does not charge by number of calls or call minutes, instead describing a unique-customer allowance (goodcall.com) (goodcall.com). A valid comparison must normalize expected activity.

Several examples illustrate the range without implying category averages:

- **Route-only flat fee:** autoattendant.io lists **\$29 monthly for the company**, not per person (autoattendant.io).
- **AI subscription plus usage:** My AI Front Desk lists **\$99 monthly**, or **\$79 monthly billed annually with 200 voice minutes** (myaifrontdesk.com) (myaifrontdesk.com).
- **AI by real call:** Smith.ai lists **25 calls monthly at \$0**, then **\$3 per call** on its Free AI tier (smith.ai).
- **Human by call:** Smith.ai lists **\$300 monthly for 30 calls** and **\$11.50 per overage call** (smith.ai).

- **Hybrid unit:** Abby lists **\$165 monthly for 50 Abby Minutes**, corresponding to up to 100 AI-answering minutes, and says a minute of AI consumes half an Abby Minute (abby.com) (abby.com) (abby.com).
- **Per-seat platform:** Dialpad's September 2026 information lists Connect Standard at **\$15 per user monthly on annual billing or \$27 monthly**, while Support Essentials is **\$80 per user monthly on annual billing or \$95 monthly** (dialpad.com) (dialpad.com).

Table 3 is a monthly-cost worksheet. Blank inputs are deliberate because call lengths, transfer behavior, annual discounts, taxes, and integration needs differ.

Cost component	Input	Monthly calculation	Verification question
Base company fee	\$_____	\$_____	Monthly or annual commitment?
Seats or receiving people	_____ at \$_____	\$_____	Are recipients, agents, or admins billable?
Included calls or minutes	_____	\$0 within allowance	What counts as a call or minute?
Overage	_____ at \$_____	\$_____	Rounded, per call, per minute, per credit, or per unique caller?
Numbers and transfer destinations	_____ at \$_____	\$_____	Are local, toll-free, and extra destinations included?
Integrations and AI credits	_____	\$_____	Are calendar, customer relationship management, and text actions extra?
Setup, onboarding, taxes, fees	_____	\$_____	One-time or recurring?
Staff cleanup	_____ hours at \$_____	\$_____	How much re-entry, correction, and follow-up remains?
Total expected monthly cost		\$_____	Recalculate at ordinary and peak volume.

The worksheet prevents a false comparison between a company-wide route-only fee, a per-seat system, a per-call receptionist, and a credit-based AI plan. It also makes hidden operational cost visible. One additional transfer destination on Smith.ai's human plans, for example, is listed at **\$15 monthly** beyond the included allowance (smith.ai). PATLive states that displayed prices exclude taxes and fees (patlive.com).

Privacy, disclosure, and retention evidence

Buyers should inventory data before comparing security badges. Ask whether the system creates audio, transcripts, summaries, caller profiles, appointment records, and payment data; where each item is stored; who receives it; how long it remains; and how it is deleted.

For U.S. healthcare organizations, HHS says a cloud provider that creates, receives, maintains, or transmits electronic protected health information on a covered entity's behalf is a business associate (hhs.gov). HHS also says storing recordings or transcripts can go beyond acting as a mere conduit (hhs.gov). HHS does not

certify particular products as HIPAA compliant ([hhs.gov](https://www.hhs.gov)).

Recording law is jurisdiction-specific. Federal law generally recognizes prior consent by one participant, subject to statutory limits ([uscode.house.gov](https://www.uscode.house.gov)). California's confidential-communication statute uses an all-parties-consent rule, subject to its exceptions ([leginfo.legislature.ca.gov](https://leginfo.ca.gov)). The Federal Trade Commission (FTC) expressly notes that state recording laws vary ([ftc.gov](https://www.ftc.gov)). There is no single nationwide rule established by these sources that every inbound AI receptionist must use the same AI disclosure. Obtain jurisdiction-specific advice.

Implications and Future Directions

The categories are converging at the feature level but not at the responsibility level. AI products add scheduling and human fallback. Answering services add automation. Phone platforms add AI reception. That makes the purchase label less reliable, so the buyer should contract for an outcome and boundary.

Five implications follow:

- **Escalation is part of the product:** Document who receives a transfer, at what hours, and what happens after no answer.
- **Knowledge needs ownership:** Assign a person to approve business facts, review unanswered questions, and update approved answers when policies change.
- **Disclosure should be explicit:** NIST recommends considering disclosure based on context, audience, purpose, risk, and frequency (nvlpubs.nist.gov). Even where a specific inbound-call announcement is not legally mandated, clear wording can set expectations and make escalation easier.
- **Accessibility needs an alternative:** A voice channel cannot be the only route for every caller. W3C advises testing with users who have varied cognitive and learning disabilities ([w3.org](https://www.w3.org)).
- **Monitoring continues after launch:** NIST recommends validating that deployed AI remains reliable in real-world scenarios ([nist.gov](https://www.nist.gov)).

For many very small businesses, the likely progression is not automatically from “basic” to “advanced.” It is from a personal phone to deliberate routing, then to conversational coverage only when the call jobs justify it, and finally to queue operations if staffing and service measurement demand them. Each step should be reversible and supported by pilot evidence.

Frequently Asked Questions (FAQs)

What is the difference between an AI receptionist and an auto attendant?

An auto attendant presents fixed choices and routes. An AI receptionist interprets natural-language requests and may answer, collect information, schedule, or transfer. The difference is conversation and action, not whether the greeting uses a synthetic voice.

What is the difference between an answering service and an auto attendant?

An answering service supplies an external person to speak with callers. An auto attendant supplies menu logic that connects callers to the buyer's own team or voicemail. If callers mainly need the right employee, routing may be sufficient. If they need judgment before anyone on the team is available, live answering may fit.

Is an AI receptionist cheaper than an answering service?

Sometimes, but not by definition. Current official examples use free allowances, subscriptions, per-call overages, credits, and hybrid human-AI units. Live services may bill per call or rounded minute. Normalize the same monthly call mix in the worksheet, including staff correction time and integrations, before comparing.

What is the best phone answering solution for a small business?

The best category is the smallest one that reliably completes the required call job. Route-only calls favor an auto attendant. Bounded intake and bookings can favor AI. Ambiguous, emotional, or high-stakes conversations can justify humans. Queues, dashboards, recordings, supervision, and multiple channels point to a full platform.

Can an AI receptionist take payments or handle medical calls?

Some products support payment-related workflows, but the buyer must verify authentication, data handling, escalation, and the exact system receiving the data. Medical use requires a separate HIPAA analysis. HHS says agreements can address use, retention, and disclosure limits ([hhs.gov](https://www.hhs.gov)).

When does a small business need a full phone platform?

Evaluate one when calls must wait for available agents, managers need abandonment or wait-time reporting, supervisors need monitoring or coaching, recordings require administrative control, or customers contact the business across voice, chat, text, email, or video. Those are operational platform requirements, not receptionist features.

How should a business choose and test a phone answering system?

List actual call jobs, rank the cost of mishandling each, identify a human fallback, normalize total monthly cost, review data flows, and run the 20-call pilot. Measure task completion, correct escalation, summary accuracy, recovery, and employee cleanup time. Re-test after scripts, integrations, or policies change.

Conclusion

AI receptionists, answering services, auto attendants, and full phone platforms solve different problems. The useful dividing line is the work required after the greeting.

An **auto attendant** is appropriate when the call should reach a known person or group and that team remains responsible for answering. An **AI receptionist** is appropriate when a repeatable conversation or action can be bounded, tested, and escalated. A **human answering service** is appropriate when nuance, emotion, or exceptions make judgment valuable. A **full platform** is appropriate when the business is operating queues, agents, recordings, dashboards, supervision, or several communication channels.

Price comparisons should reflect the underlying unit, including company fee, seats, calls, minutes, credits, overages, integrations, taxes, and cleanup labor. Privacy review should cover audio, transcripts, summaries, retention, subprocessors, consent, and healthcare obligations where applicable.

The final decision should come from observed calls. A small, deliberately difficult pilot can reveal dead-end transfers, incorrect actions, inaccessible flows, and avoidable staff work before those failures reach customers. There is no universally best category. There is a best-fit operating model for the call jobs, risk level, staff availability, and evidence each business can verify.

Source links

Links cited in this article, in order of first appearance. Full addresses are provided for readers using extracted text.

1. <https://learn.microsoft.com/en-ca/MicrosoftTeams/setting-up-your-phone-system>
2. <https://docs.aws.amazon.com/connect/latest/adminguide/about-routing.html>
3. <https://learn.microsoft.com/en-us/microsoftteams/manage-queues-app>
4. <https://autoattendant.io>
5. <https://smith.ai/pricing/receptionists>
6. <https://www.myaifrontdesk.com/pricing>
7. <https://airc.nist.gov/airmf-resources/airmf/3-sec-characteristics/>
8. <https://www.hhs.gov/hipaa/for-professionals/special-topics/health-information-technology/cloud-computing/index.html>
9. <https://learn.microsoft.com/en-us/microsoftteams/aa-cq-setup-auto-attendant-cmdlets>
10. <https://www.ringcentral.com/air>
11. <https://www.ruby.com/answering-service/>
12. <https://docs.aws.amazon.com/connect/latest/adminguide/contact-events.html>
13. <https://www.w3.org/WAI/WCAG2/supplemental/patterns/o6p04-voice-menus/>
14. <https://assets.ringcentral.com/us/datasheet/rc-ai-receptionist-privacy-datasheet.pdf>
15. <https://www.ringcentral.com/legal/dpa-subprocessor-list.html>
16. <https://www.pnas.org/doi/full/10.1073/pnas.1915768117>
17. <https://smith.ai/ai-receptionist>
18. <https://smith.ai/features/appointment-scheduling-service>
19. <https://www.answerconnect.com/services>
20. https://docs.aws.amazon.com/us_en/connect/latest/adminguide/amazon-connect-contact-centers.html
21. <https://www.genesys.com/capabilities/wem-workforce-engagement-management>
22. <https://media.zoom.com/download/assets/zoom-contact-center-privacy-data-sheet.pdf/e09a9c04104d11eeb7b04eb1d1916cf2?lang=en-US&onlycontent=1>
23. <https://airc.nist.gov/airmf-resources/airmf/5-sec-core/>
24. <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf>
25. <https://www.w3.org/WAI/people-use-web/abilities-barriers/speech/>
26. <https://www.nist.gov/publications/effects-speech-recognition-accuracy-performance-darpa-communicator-spoken-dialogue>

27. <https://www.goodcall.com/pricing>
28. <https://smith.ai/pricing/ai-receptionist>
29. <https://www.abby.com/products/ai-receptionist/>
30. <https://www.dialpad.com/llm-info/>
31. <https://www.patlive.com/pricing/>
32. <https://www.hhs.gov/hipaa/for-professionals/privacy/guidance/hipaa-audio-telehealth/index.html>
33. <https://uscode.house.gov/view.xhtml?req=%28title%3A18+section%3A2511%29>
34. https://leginfo.ca.gov/faces/codes_displaySection.xhtml?lawCode=PEN§ionNum=632
35. <https://www.ftc.gov/business-guidance/resources/complying-telemarketing-sales-rule>
36. <https://www.nist.gov/publications/challenges-monitoring-deployed-ai-systems-center-ai-standards-and-innovation>

THE PUBLISHER

About Auto Attendant

Auto Attendant gives a small business one phone number and a recorded menu that sends each caller to the right person. The call rings the mobile that person already carries, using their normal dialler and their normal mobile plan. Pricing is a single flat monthly fee for the whole company rather than a charge for each user, and there is no app, desk phone or hardware to install.

One number, one menu, the phones you already own

A caller dials the business number, hears a short greeting and chooses an option. Each option can ring one mobile, ring several at once so whoever is free answers first, or try people one after another in a set order. There is no limit on how many people sit on the receiving end, because nobody is billed per seat. Setup is done by the owner in a few minutes: pick a number, write the greeting, and point each option at a mobile.

What it deliberately is not

An auto attendant routes a call; it does not answer one. Nobody at Auto Attendant speaks to callers, and no AI stands in for a receptionist. A business that wants its calls handled, messages taken and appointments booked wants an answering service, which is a different product at a different price. A business that needs call recording, queues, agent dashboards or CRM integration wants a contact centre platform. Auto Attendant is for the case where the call should simply reach the right pocket.

Keeping personal numbers off the internet

The business number is the only number a caller ever sees. Personal mobile numbers are never displayed and never given out, which is what lets an owner put a number on a website, an invoice, a van or a Google Business Profile without handing a personal line to everyone who finds it. An existing advertised number can be ported across and kept, and the old line keeps working while the port is in progress.

Work with Auto Attendant

Visit [Auto Attendant](#) to see [how it works](#), [how it compares](#) to a personal mobile, Google Voice and a full phone platform, and the current [pricing](#).

Telephone number portability, emergency calling and caller-ID rules differ by country and change over time. Articles published here describe general practice and cite primary regulator and carrier sources; they are not a guarantee of any timeline, and they are not legal or regulatory advice for a specific business.

Website: <https://autoattendant.io>

This article is educational. Verify consequential medical, legal, financial, regulatory and technical decisions with appropriate professionals and the cited primary sources. Company services are described separately from the article's research findings.