




FIVE CHECKS

Five Checks Before You Add an AI Receptionist

This guide gives five checks to run before launch. The goal: avoid duplicate work, missed follow-up, and staff who are busier than before.



David Cyrus

Attainment

Your front desk answers calls, books patients, moves appointments, handles urgent concerns, and follows up on messages. During a busy check-in, some calls go to voicemail and some never get a callback. An AI receptionist can help with part of this. Added without a plan, it can create a second pile of work.

This guide gives five checks to run before launch. The goal: avoid duplicate work, missed follow-up, and staff who are busier than before.

Key takeaways

- Start with the problem, not the tool. An AI receptionist is not right for every practice.
- Decide call by call what the AI may answer, capture, route, or support.
- A pleasant conversation is not a booking. Verify each practice-software action before promising it to patients.
- Every request the AI cannot finish needs one named owner and one clear next step.
- Test real calls before launch and review the whole chain, not just the recording.

THE FIVE CHECKS

01

Map calls and requests



02

Decide the AI's role per request



03

Verify every practice-software action



04

Set the handoff when the AI stops



05

Test real situations and review the chain

01

Check 1: Map the calls and requests your front desk handles today

Direct answer: Before you add anything, write down what comes in and what your team does with it.


New-patient requests, reschedules, cancellations, insurance and billing questions, recalls, and urgent dental concerns all land on the same line.

Questions to ask:

- What problem are we trying to solve? Missed calls, slow callbacks, after-hours coverage, or front-desk overload?
- Which calls go to voicemail, and how many get a callback the same day?
- Which call types take the most staff time?

Example: An office pulls two weeks of call logs and voicemails. Reschedules and new-patient inquiries make up most after-hours messages. Urgent concerns are rare but always need a person. That shows where the AI might help and where it must not act alone.

Output: A one-page list of request types, how often each happens, who handles it today, and what good looks like.



02

Check 2: Decide which requests the AI may answer, capture, route, or support

Direct answer: Sort each request type into one of four roles: answer, capture, route, or support.

Answering means giving information such as hours or parking. Capturing means collecting details for staff to act on. Routing means sending the request to the right person. Supporting means helping the caller while staff decide.

Questions to ask:

- Which calls should the AI handle, and which must always reach a person?
- What information may it provide, and what may it collect?
- Which requests need clinical or staff judgment before any action?
- What should the AI say when a request is outside its role?

Example: A new patient calls to book a first visit. The AI may capture name, phone number, reason for visit, and preferred days. It may not confirm a time until Check 3 verifies that action, so the caller is told staff will confirm.

Example: A patient calls with a swollen jaw and pain. In office hours the AI captures the concern and routes it to staff right away. After hours it gives the approved emergency instruction and records the message for staff the next morning. It never judges how serious the problem is.

Output: A table with each request type, its role, what the AI may say or collect, and what stays with staff.

SORT EACH REQUEST TYPE INTO ONE OF FOUR ROLES



Answer

giving information



Capture

collecting details



Route

sending the request
to the right person



Support

helping the caller
while staff decide

CLINICAL OR STAFF JUDGMENT

a request is outside its role

stays with staff

03

Check 3: Confirm every practice-software action separately before promising booking or record updates

Direct answer: Do not tell patients the AI can book, move, or update anything until you have seen it work in your own practice software.

A vendor may say the AI connects to your practice software. That can mean it reads the schedule, writes to it, or only sends a note to an inbox. Each action is separate and must be tested in your live system, not a demo.

Questions to ask:

- Can it actually read the schedule, or only capture a request?
- Can it write a new appointment, or does staff enter it?
- Can it change an existing appointment without creating a duplicate?
- Where is the request recorded if the action fails?
- Is our system and version supported, and verified with our data?

Example: An existing patient asks to move a cleaning from Tuesday to Thursday. The conversation sounds perfect, and the patient is told it is done. When staff check, Tuesday is still booked and no Thursday slot exists. The conversation succeeded. The workflow failed.

Output: Every expected software action marked verified, captured only, or not supported. Nothing is promised to patients until marked verified.

EACH ACTION IS SEPARATE AND MUST BE TESTED IN YOUR LIVE SYSTEM

read schedule

book

change

cancel

update record

Every expected software action marked verified, captured only, or not supported.

04

Check 4: Give staff one clear next step when the AI cannot finish a request

Direct answer: Every request the AI cannot complete needs a named owner, a landing place, and a response time.

The AI will stop sometimes. It will misunderstand a caller, hit a request outside its role, or lose a connected system. That is fine if staff know what happens next. It fails when a request sits in a queue nobody checks.


Questions to ask:

- Who owns the next step when the AI stops?
- Where does the request land: a task list, inbox, text, or software note?
- How quickly must staff respond, during hours and after hours?
- What is the fallback when a system or the AI provider is unavailable?
- What does the caller hear when the AI hands off?

Example: A caller asks about a procedure the AI is not set up to discuss. The AI says so, takes the caller's name and number, and creates a callback request on the front-desk lead's task list. The rule is a callback within two business hours. If the lead is away, the office manager covers.

Output: A written handoff rule for each request type: owner, landing place, response time, backup owner, and fallback when systems are down.

WHEN THE AI STOPS



Caller told what happens next at every handoff.

05

Check 5: Test real call situations and review the full chain before expanding

Direct answer: Run realistic test calls and check six things for each: conversation, system action, routing result, staff notification, privacy behavior, and audit record.


A recording is not enough. Follow each request to the end and confirm every link worked.

Questions to ask:

- How will we know the full workflow worked, not just the conversation?
- Did the right staff member get notified in time?
- Was any information collected or shared that should not have been?
- Is there a record showing what the AI said and did?

Example: A test caller leaves a message at 9 p.m. asking to reschedule. The next morning, did the request appear where the front desk looks first? Was it marked for action? Was the record complete? If any part is missing, the test fails, even if the conversation was smooth.

Output: A test log for each situation: new-patient booking, appointment change, urgent concern, a request the AI does not understand, and an after-hours message. Each row shows pass or fail on all six.



What good handoff evidence looks like

Good evidence is something a staff member can point to, not a vendor dashboard or a vendor's word. For each test call, show:

- ✓ The conversation record: what the caller asked and the AI said.
- ✓ The system action: a screenshot or record from your practice software, or a note that nothing changed.
- ✓ The routing result: which person or queue received the request.
- ✓ The staff notification: a timestamp showing the right person saw it.
- ✓ The privacy behavior: only the agreed information was collected and shared.
- ✓ The audit record: a trace of the request if a patient asks later.

If any one is missing, you have evidence that a conversation happened, not that the workflow worked.

CONVERSATION

A conversation happened.

NOT

FULL CHAIN

The workflow worked.

When to pause the launch

Hold the launch if any of these are true:

- A request type has no named owner for the next step.
- A software action is promised to patients but not verified in your live system.
- There is no fallback when the AI or a connected system is unavailable.
- A critical test failed. Urgent concerns, after-hours messages, and appointment changes are always critical.
- Staff do not know where handed-off requests land or how fast to respond.
- You cannot produce an audit record for a test call.
- Privacy and consent requirements have not been reviewed against what the AI collects and stores. Confirm your own obligations and vendor agreements. This guide is not legal advice.

A hold is not a failure. It is cheaper than cleaning up double bookings and missed patients.

FAQ

Does every dental practice need an AI receptionist?

No. If your calls are answered and followed up the same day, you may not need one. Start with Check 1.

Will it connect to our practice-management software?

Maybe. Not every AI receptionist works with every system, and support is not a verified action. Confirm each action separately.

Can it replace a front-desk staff member?

Plan for it to reduce some interruptions and capture some requests, not to replace staff. Every check depends on staff owning the next step.

How do we know a booking actually happened?

Only when the system action and final record in your software confirm it. The AI saying "you are booked" is not confirmation.


What about privacy?

Know what the AI collects, where it is stored, who can see it, and how long it is kept. Confirm this against your own requirements and vendor agreements.

Closing

AI automation can take real weight off a busy front desk when the role is clear, software actions are verified, and staff always know the next step. It can also add work when those pieces are missing. Use the one-page checklist to run the five checks with your team.

THE FIVE CHECKS



Launch decision



Ready

All sections complete.
Critical tests passed.
Evidence on file.



Hold

Any item open. Fix and retest before launch.



Not ready

Problem, AI role, or software actions not defined. Return to Check 1.

Front-Desk AI Readiness Checklist

Companion to "Five Checks Before You Add an AI Receptionist." Check a box only when you can show it is done.

Practice: _____

Date: _____

Completed by: _____

01 Map calls and requests

- ☐ Reviewed two weeks of call logs and voicemails for frequency and who handles each.
- ☐ Listed request types: new patient, reschedule, cancellation, insurance, billing, recall, urgent concern, other.
- ☐ Wrote down the one problem to solve.

02 Decide the AI's role per request

- ☐ Each request type marked: answer, capture, route, or support.
- ☐ Listed what the AI may say and collect.
- ☐ Urgent concerns always reach staff judgment; the AI never rates severity.
- ☐ Wrote what the AI says to out-of-role requests.

03 Verify every practice-software action

- ☐ Listed expected actions: read schedule, book, change, cancel, update record.
- ☐ Each action marked: verified in our live system, capture only, or unsupported.
- ☐ Tested for duplicates when changing a booking.
- ☐ Nothing promised to patients until marked verified.

04 Set the handoff when the AI stops

- ☐ Every request type has an owner, backup owner, and landing place.
- ☐ Response times set for office and after hours.
- ☐ Fallback set for AI or connected system outages.
- ☐ Caller told what happens next at every handoff.

05 Test real situations and review the chain

For each test, confirm all six: conversation, system action, routing, staff notification, privacy behavior, audit record.

- ☐ New-patient booking: pass/fail
- ☐ Appointment change: pass/fail
- ☐ Urgent concern: pass/fail
- ☐ Request the AI cannot understand: pass/fail
- ☐ After-hours message seen next morning: pass/fail
- ☐ Evidence on file for each test, not just a recording.
- ☐ Privacy and consent reviewed against what the AI collects and stores, including our obligations and vendor agreements.

Launch decision

Rule: Any missing owner, unverified system action, missing fallback, or failed critical test means Hold.

☐ **Ready:** All sections complete. Critical tests passed. Evidence on file.

☐ **Hold:** Any item open. Fix and retest before launch.

☐ **Not ready:** Problem, AI role, or software actions not defined. Return to Check 1.

Decision: _____

Signed: _____

Review date: _____