AGENT READINESS ASSESSMENT

The six dimensions

Before we write a line of agent code for anyone, we score their sales system on six things. This is that instrument written out in full: what each dimension measures, why it matters, what good looks like, and what to fix first.

Take the assessment

Download as PDF

How to use this

It's meant to be used, not just read. Run it on your own organization in a room with the people who would own the result. You don't need us there.

1. **Place your organization on each row of the matrix.**

   Be honest rather than aspirational. The value of the exercise is entirely in where it disagrees with the story you tell externally.

2. **Expect an uneven profile.**

   Nearly every organization is agent-ready on two or three dimensions and absent on one. The absent one is the finding, and it is usually the one nobody owns.

3. **Sequence the work. Do not fix everything at once.**

   The six are not equally weighted and they are not independent. The reading order is at the bottom of this page.

The same six dimensions with the site chrome stripped, for circulating internally.

The maturity matrix

Six dimensions, four levels. Most organizations land in more than one column. That's the point: the shape of the profile tells you where to start.

Agent readiness maturity by dimension. Go Fig, 2026.

Dimension	Absent No mechanism at all	Manual Happens only because a person holds it together	Assisted Partly systematic, real gaps remain	Agent-ready Automated, observable, and trusted
01 Systems connectedness	Scattered and nothing talks to anything	Many, stitched together by hand	Several, partly connected	Connected, syncing automatically
02 Admin load	Admin eats the week and selling happens in the gaps	Mostly admin and follow-up, roughly 20/80	Roughly balanced, 50/50	Mostly selling, roughly 80/20
03 CRM hygiene	The forecast is a fiction rebuilt by hand	Reps clean it up right before the pipeline review	Mostly current, some deals go stale	Current and every open deal has a real next step
04 Follow-up latency	Often never and leads go cold on their own	A few days if someone remembers	Same day	Within minutes because it is automated
05 Signal response	Nothing until we read the press release, if ever	Someone notices eventually but usually too late	A rep sees it and acts that week	Automated and the account surfaces with context attached
06 Traceability and approval	No way to tell so we do not automate at all	We would see the result, never the reasoning	We could reconstruct most of it with some digging	Every action logged and approved before it sends

The dimensions in full

01

Systems connectedness

How many systems hold your sales data and do they reconcile without a person?

WHAT IT MEASURES

How many systems hold a fact about an account, and whether those facts reconcile without someone doing it. Count the CRM, email, calendar, call recordings, support tickets, product usage, billing, and the spreadsheets in between.

WHY IT MATTERS

An agent can only reason over what it can see. Point one at the CRM alone and it inherits every one of the CRM's blind spots. It will draft a confident follow-up to an account that filed three support tickets last week because nothing told it otherwise. The most common cause of an agent that looks unreliable in production is not the model. It is an agent reasoning over a partial record and being right about the wrong thing.

WHAT GOOD LOOKS LIKE

The CRM, email, calls, tickets, product usage, and billing land in one place on a schedule you control. That place is infrastructure you own rather than a vendor's copy of your business.

WHAT TO FIX FIRST

Inventory before integration. List every system that holds a fact about an account and mark which ones are able to share data with other tools (an API) today. The gap list is your integration backlog. It's almost always shorter than the room expects.

02

Admin load

How much of a rep's week goes to CRM admin and manual follow-up versus actually selling?

WHAT IT MEASURES

The split between administrative work (logging calls, updating stages, writing recaps, chasing next steps) and live conversations with buyers. Measured per rep, per week.

WHY IT MATTERS

This is the dimension that funds the program. Hours lost to admin turn straight into a dollar figure your sales lead can take to the budget owner. If you can't state the number, you can't make the budget case. The project then stalls at the budget conversation rather than the technical one.

WHAT GOOD LOOKS LIKE

The post-call work is drafted before the rep reopens their laptop. Nobody is retyping into a CRM something they already said out loud on a call.

WHAT TO FIX FIRST

Measure it before you automate it. Take a two-week sample across five reps and count honestly. The baseline is what your pilot gets judged against. A pilot without a baseline gets judged on vibes.

03

CRM hygiene

How current and trustworthy is your CRM?

WHAT IT MEASURES

Whether open deals carry a real next step and an honest close date, continuously, rather than in the hours before a pipeline review.

WHY IT MATTERS

Agents act on what the data claims. A close date three weeks in the past isn't a neutral blank. It's an instruction and an agent will follow it. An agent running on a dirty CRM nudges the wrong deals, skips the right ones, and burns the team's confidence inside two weeks. That trust gets spent once. Recovering it costs more than the cleanup would have.

WHAT GOOD LOOKS LIKE

Hygiene is enforced continuously by something that never forgets, instead of reconstructed by people under deadline pressure.

WHAT TO FIX FIRST

This is the strongest candidate for a first agent, because it is the rare one that improves the data it runs on. It's a nightly check for deals with no next step, no activity in two weeks, or a close date already past. It posts the list to the owner. Once the rule is approved, it fixes the fields itself.

04

Follow-up latency

After a call or an inbound signal, how long until follow-up actually happens?

WHAT IT MEASURES

Elapsed time from a triggering moment to the thing that follows it: the recap email, the booked next step, the reply to a demo request, or a pricing-page visit.

WHY IT MATTERS

This is where agents produce revenue rather than savings. Every other dimension gives time back. This one saves deals that would otherwise be lost to silence. That's a much better line on a business case. It also has the cleanest before-and-after measurement of the six, which makes it the pilot most likely to survive scrutiny.

WHAT GOOD LOOKS LIKE

Same hour every time. The message is drafted automatically and a person sends it until the team trusts it enough to stop reading every one.

WHAT TO FIX FIRST

Measure it. Most organizations have never measured median time-to-follow-up and are genuinely surprised by the number. The measurement alone changes behavior before a single agent ships, which makes it the cheapest intervention on this page.

05

Signal response

When a target account raises a round, hires an exec, or launches something, what actually happens?

WHAT IT MEASURES

Whether externally visible buying signals at your target accounts reach the person who could act on them, while they still matter.

WHY IT MATTERS

Nobody fails this one on purpose. It fails because it has no owner. Work without an owner never reaches a roadmap. It is also the dimension best suited to automation on first principles: it is pure sustained watching, which is the thing humans are worst at and machines are indifferent to.

WHAT GOOD LOOKS LIKE

The account surfaces on its own with the hook attached and the outreach already drafted. A rep only has to approve it, not discover it.

WHAT TO FIX FIRST

Write the trigger inventory. Which events would genuinely change a rep's priorities this week? Five is plenty. A signal agent watching five well-chosen events beats one watching fifty, because the second one gets muted by week three.

06

Traceability and approval

If an automated action went out on your behalf, could you tell what it did, what data it read, and why?

WHAT IT MEASURES

Whether you can reconstruct what an automated action did, which records it read to decide, and who approved it. And whether a person could have stopped it first.

WHY IT MATTERS

This is the dimension that decides whether the program ships at all. The other five determine whether agents will work. This one determines whether you will be permitted to run them. Security review, legal, and the executive who owns the customer relationship all ask the same question, and there is no version of the answer that starts with we would see the result. Programs that defer this build something impressive in pilot and then die at the enterprise readiness review, having spent the political capital and produced nothing deployable.

WHAT GOOD LOOKS LIKE

Every read, action, and approval is logged. Whether an action runs on its own or waits for approval is set per agent and per action by the business owner, and changing it takes a click, not a project.

WHAT TO FIX FIRST

Decide on paper, before the first agent, when an action gets to run on its own. Everything starts out waiting for your OK. The agent drafts, a person clicks send. Then define what earns an action the right to run on its own, whether that is a volume of clean approvals, a rejection rate under a threshold, or a named owner's sign-off. Writing this down early is free. Retrofitting it is not.

Where to start when your profile is uneven

The six dimensions do different jobs. Treating them as a checklist to finish in order is the most common way this exercise gets wasted.

Want the score instead?

The assessment turns the six dimensions into six questions. In about two minutes you get your score, the hours and pipeline risk behind it, and what to fix first.

See which of these Go Fig would run for you first

A solutions engineer walks through your pipeline with you and maps which admin the agents can take off your team. You approve every send.

Schedule a demo

Go Fig

The sales admin for sales reps. Greenville, SC.

support@gofig.ai · (828) 513-0465

Sales admin notes by email

name@company.com Subscribe

Short notes on selling with less admin. Unsubscribe anytime.