



— A SOCRATIX POINT OF VIEW

UI for AI, UI for Agents

Is your business ready to talk to a lot of agents?

HUMAN WISDOM. AI ILLUMINATION.

Socratix AI • Cleveland, Ohio • October 2026
socratixai.com • info@socratixai.com

"User interface" assumed a human on one side and a program on the other. That just broke in two places.

For thirty years "user interface" meant one thing: how a person uses software. Screens and buttons and forms. The whole discipline assumed a human on one side and a program on the other, and that assumption just broke in two places at the same time.

The software side got a mind of its own, for one. A person working with an AI agent isn't filling out a form, they're delegating and reviewing and correcting, and the screens that were built for data entry are the wrong screens for that.

And the human side is increasingly not a human. When your sales inbox gets an email from a customer's purchasing agent, when your AP system receives an invoice that a vendor's agent submitted, when one of your own agents hands work to another one, the "user" of your business is a piece of software. It needs an interface too, and almost nobody has built it one.

— UI FOR AI

The chat box is the wrong interface for most work.

The default AI interface right now is a blank text box with a blinking cursor. It's fine for exploring and poor for operations. It asks the person to know what to ask, type it well, read a wall of text back, and remember to come back and do it again tomorrow, and most people won't. RSM's 2026 survey found 85% of middle-market executives say leadership is more enthusiastic about AI than employees are, and some of that gap is the box. Leaders saw the demo. Employees got a cursor.

The interfaces that work put the AI where the work already happens and ask people to do something they already know how to do, which is approve, edit or reject.

The big platforms have noticed. At Dreamforce in September Salesforce launched a layer meant to bring its data and permissions into whatever interface people and agents happen to be working in, and Marc Benioff called what's happening "an interface revolution." Gartner has projected that 40% of enterprise applications will include task-specific agents by the end of 2026, up from under 5% last year.

So there are two design problems now where there used to be one. How a person and an agent work together, and how your business presents itself to software that wants to do business with it. Most mid-market companies are working on neither. The second one is the one that decides who gets the order.

One agent we built for a client is set to run once a month. It reads employee calendars, works out the vacation each person took, and sends each of them a card in Teams: here's what I found, confirm or edit. The manager gets the confirmed list and approves it or sends it back, and the record is written. Nobody prompts anything. Nobody "uses AI." They approve a card, and accounting can still open the record and change it by hand, which is exactly right at this stage.

What a good human-AI interface does, roughly in the order we see it missing. It shows up in the flow of work. Teams, email, the queue somebody already watches. Not a new destination people have to remember to visit.

It shows its work: what it read, what it inferred, what it's unsure about. Salesforce's own agent design guidance makes the same point, that AI isn't perfect and users know it. Hide the reasoning and people stop trusting the output. Show it and they start correcting it, which is what you want.

It makes correction one click, because the edit path has to be shorter than the workaround or people take the workaround. It keeps the human decision visible, who approved what and when. That's your audit trail and your training signal in one place. And it doesn't ask people to prompt. The prompt is the agent's job. The person's job is judgment.

— UI FOR AGENTS

Your business as something software can use.

Now turn it around. The thing that needs to use your business is an agent. What does it need?

Not screens. An agent working through a human interface, taking screenshots and clicking the way a person would (the industry calls this "computer use"), is slow and fragile and expensive. It works as a bridge to systems nobody has opened up yet. It isn't a way to run a business.

It needs a way to read, first of all. Structured data behind an endpoint, or at minimum documents in a place it can reach with permission. The PDF report that gets emailed every Friday is readable. The same report printed and put in a binder isn't. Whether the data has to be clean before any of this works is the subject of *The Future of Data and Organizations in the AI Era*. Short answer, no.

It needs to know what it's allowed to do, which means permissions scoped to the agent, by role, the same way you'd scope a new hire. A common question: how do we keep the sales team's agents out of the M&A files? Same way you keep the sales team out. Put that data in its own place in your tenant with its own access. Full admin control and no new policy to write.

This is a behavior problem before it's a software problem. People adopt what fits a habit they already have. The card in Teams works because people already approve things in Teams. A new portal with a new login, however good it is, is asking for a new habit, and new habits are what most AI rollouts quietly die of.

It needs a way to write back that the system of record will accept. A well-formed record, through the same door a person would use, with the same approval.

And it needs a shared language, which until recently didn't exist. Two open standards have shown up. MCP, the Model Context Protocol that Anthropic released in late 2024, standardizes how an agent reaches tools and data. A2A, Agent2Agent, released by Google in 2025 and now under the Linux Foundation, standardizes how agents find each other and work together. By spring of this year more than 150 organizations were backing A2A, Microsoft and SAP and Salesforce and ServiceNow among them. One analyst's shorthand holds up pretty well: MCP is how an agent reaches things, A2A is how agents reach each other.

What this means for a mid-market company is simpler than it sounds. You don't need to pick a protocol. Your Microsoft tenant and your ERP vendor and your bank will pick for you, and they're converging anyway. What you need to do is stop building things only a person can read. Every PDF report, every "just email me the spreadsheet," every process that lives in one person's head is a door an agent can't walk through, and the job for the next couple of years is opening those doors in priority order, with the right permissions on each one.

1 · A way to read

An endpoint, or documents it can reach with permission.

PDF EMAILED FRIDAY ✓ · PRINTED AND FILED ✕

2 · A way to know what it may do

Permissions scoped per agent, by role, like a new hire.

SALES AGENTS NEVER SEE THE M&A FOLDER

3 · A way to write back

A well-formed record, through the same door, with the same approval.

THE ERP ACCEPTS IT LIKE ANY ENTRY

4 · A shared language

MCP to reach things, A2A to reach each other.

YOUR VENDORS ARE CONVERGING ON BOTH

— THE TEST

Is your business ready to talk to a lot of agents?

Sometime in the next year or two a customer's purchasing agent is going to email your sales inbox asking for a quote on 500 units with the spec attached, and it will take a structured answer in minutes or move on to the next supplier. A vendor's agent will submit an invoice and ask when it's getting paid. Your bank's agent will ask for three months of statements to renew the line. A candidate's agent will ask about your benefits before the person decides whether to apply.

Can your business answer, or does each one land in a queue behind somebody who's out until Thursday?

The inside version is just as real. Your own agents will need to talk to each other. One reads the work order and hands off to one that checks inventory, which hands off to a person who approves the exception. If each of those was built as a one-off chatbot they can't, and if each was built with a readable input, a scoped permission and a defined write-back they can, and the third one takes a week instead of a quarter.

There's a reason the industry is already talking about moving from paying for tokens to paying for outcomes. When the thing you're buying is a finished process, the business that can hand an agent a legible process gets the result first and cheapest. I don't think the companies that come out ahead will be the ones with the most AI. They'll be the ones whose processes were clear enough for any agent to use. How to measure the result so it survives a budget review is in *Measuring What Matters*.

What we build to. They hold whether the next user of a process is an agent or a new hire.

01

One system of record, and everything else is a view.

The ERP, the ledger and the HR system own the truth. Agents read from them and write back to them and never keep a private copy.

02

The agent proposes and a person approves,

until the measured error rate earns it more. Autonomy gets granted one step at a time, with numbers.

03

Every output shows its sources.

What was read, what was inferred, what's uncertain. For a person that's trust; for the next agent in line it's the input.

04

Permissions by role, scoped per agent.

An agent gets what the person it works for would get. Sensitive data sits in its own place with its own access.

05

Build for the second agent.

If this one's output could be another one's input, structure it that way now. A table beats a paragraph and a field beats a table.

06

Same request, same result.

Moving a vacation day from the 3rd to the 4th is one request, and running the process twice shouldn't create two records. Engineers call this idempotency. Everyone else calls it not double-counting.

07

Package and version what you build.

A test environment before production, a managed solution before a broad rollout, and a record of what changed, because IT will ask and because the second version shows up faster than you expect.

— WHERE TO START

The interface inventory. Five processes, three columns.

Take your five most important processes and for each one fill in three columns.

PROCESS	HOW A PERSON DOES IT TODAY	HOW AN AGENT COULD READ IT	HOW AN AGENT COULD WRITE BACK
Quote request	<i>Email arrives, rep opens the spec PDF, prices it in a spreadsheet, replies</i>	<i>Shared inbox and spec attachment, with the pricing sheet readable</i>	<i>Draft reply for the rep to approve; quote logged in the CRM</i>
Vendor invoice	<i>AP clerk re-keys the PDF into the ERP, matches to PO by hand</i>	<i>Invoice PDF and open POs, both reachable</i>	<i>Matched entry proposed for AP approval; written to the ERP</i>

Start where a person is already doing the translating. The one who turns a spreadsheet into an answer, or re-keys the customer's PO into the ERP. That person is the interface today. Build the agent to do the translating, keep the person on the approval, and the first door is open.

Then the next one. It climbs the way it always does, a personal agent for one person, then a use case for a team, then something shared across departments, then enterprise-wide, and each process you open makes the next one cheaper because the permissions and the pattern and the write-back are already sitting there.

— KEEP READING

The Future of Data and Organizations in the AI Era

Your data is a mess, and that stopped being the problem. The other half of this paper, written from the data's side.
socratixai.com

Measuring What Matters

A foundations-based approach to AI ROI. The bubble won't punish spending. It will punish unmeasurable spending.
socratixai.com

The Three-Foundation Framework

How mid-market companies turn AI from a curiosity into compounding value. The model behind everything we write.
socratixai.com

About Socratix AI

Socratix AI is a Cleveland firm that turns AI ambition into completed work for mid-market companies: strategy, training and hands-on implementation. We work across three foundations, People & Culture, Business Strategy, and Capability & Technology, because an interface is where the technology meets the people who have to live with it.

If you want to know how many of your processes an agent could use today, we'll walk through them with you. It starts with a conversation and a walk through the operation. info@socratixai.com · socratixai.com

Human Wisdom. AI Illumination.

SOURCES Salesforce, AIforce press release (September 16, 2026) · Salesforce blog, Get 5 Insider Design Tips for Building Your AI Agents · RSM Middle Market AI Survey 2026: U.S. and Canada (RSM US, July 2026) · MintMCP, A2A vs MCP: What Each Protocol Does (September 2026) · Beam, Agent2Agent vs MCP: The 2 Protocols Your 2026 AI Agent Stack Runs On (citing Gartner, August 2025).