

# Where Systems of Intelligence Begin

Accounts Receivable Goes First

AUTHOR:



Emma Rowand

Research Analyst



### About Activant

Activant is a research-led global investment firm that partners with high-growth companies. Since 2015, we have invested in category-defining businesses during their most critical phases of growth, partnering with founders who have won their initial battles and are ready for the next challenge.

Our approach pairs deep, proprietary research with patient, flexible capital and hands-on operational partnership. We work alongside founders and leadership teams to refine strategy, strengthen operations, and accelerate sustainable growth.

Activant Research is dedicated to uncovering the most exciting emerging technologies, sectors, and companies we believe will shape the future. Our research-driven perspective informs everything we do, helping us invest at meaningful inflection points and support founders in building enduring, category-leading businesses.

You can find out more about Activant and our research at <https://activantcapital.com/>.

Last week we published [AR Operations Today: Satisfied But Stuck—And Ready For Agents](#), drawing on a proprietary survey of 430 accounts receivables (AR) leaders.

The headline finding: **48% of companies with more than a billion dollars in annual revenue still run AR on spreadsheets.** These are the same companies that have implemented SAP, hired CFOs with measurable operating budgets, and spent the last decade announcing digital transformation initiatives. Their AR clerks are still copy-pasting remittance data from PDF email attachments into pivot tables.

That isn't because the people running the function are resistant. **90%** of AR leaders are comfortable with AI agents handling tasks autonomously or with light oversight. **96%** find role-tiered agent pricing appealing. And **54%** say faster implementation and time to value is the single thing most likely to make them switch tools in the next six months.

So we have a function that's broadly automatable, run by people who want to automate it, with buyers eager to pay for it, which somehow remains stubbornly manual at scale.

## The Spreadsheet Tell

This isn't a story about laggards or under-investment. It's a story about what Systems of Record (SORs) were never built to do.

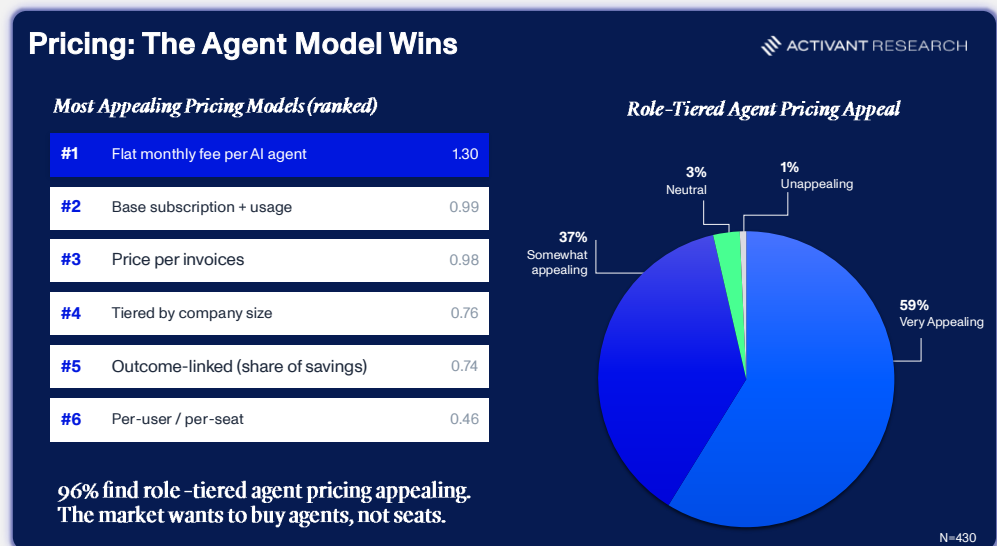
The visible AR software TAM sits around \$4B.<sup>1</sup> The top five vendors (SAP, Oracle, SK Global Software, Quadient, and Workday) control just under 40% of global revenue.<sup>2</sup> Four were founded before the dot-com boom. Quadient dates back to 1924.

These systems earned their place by digitizing the paper trail, replacing filing cabinets with databases. But their schemas reflect the constraints of their era. Every field, dropdown, and form was shaped around what a human could realistically type into a box. That's why clerks keep retreating to Excel. When a customer short-pays an invoice because half the order shipped late and the other half had a pricing dispute carried over from last quarter, none of that fits the SOR's fixed menu of reason codes. So the clerk opens a spreadsheet, invents her own codes, writes her own context, and gets the work done herself.

A quarter of respondents named high exception volume requiring human judgment as the biggest driver of inefficiency. Another 21% pointed to manual, repetitive work. **That's 46% of the market saying the real problem is everything the software can't handle.**

A System of Intelligence (SOI) responds differently. It doesn't try to enumerate exceptions. It reads the remittance email, parses the attached PDF, joins it against the open invoice ledger, makes a decision, and escalates only what it's genuinely uncertain about. The queue shrinks because the system is reasoning about the exception rather than routing around it. Activant portfolio company, [Stuut Technologies](#) is doing exactly that, autonomously handling collections, cash application, and disputes end-to-end, live on top of your existing finance stack.

Once a system handles that work, the buying motion changes, and buyers are the ones changing it. When we asked AR leaders to rank six pricing models, a flat monthly fee per AI agent ranked first in every revenue tier. Per-seat licensing ranked last. The 96% who like role-tiered pricing want different agents priced at different levels based on the complexity of the work, the way you'd pay a junior clerk less than a senior collections analyst. **Buyers have stopped benchmarking AI against software budgets and started benchmarking it against payroll.**

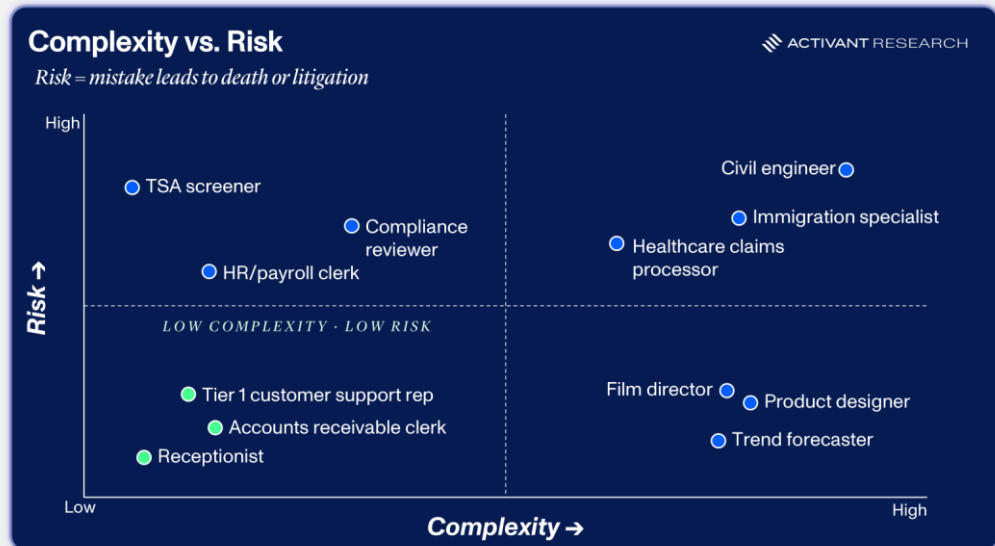


That changes the TAM entirely. There are over 676,000 AR clerks in the United States earning \$44,000 to \$58,000 a year.<sup>3,4</sup> At the low end, that's roughly **\$29.8B in annual labor costs**, or **7.4x global AR software spend**. And that's clerks alone, before the analysts, specialists, managers, and directors above them. When the buyer's mental model shifts from "what does the tool cost" to "what does the work cost," the ceiling on enterprise spend moves from the IT line to the payroll line.

SORs were priced per seat because value scaled with the number of humans using them. SOIs do the work themselves. Pricing them per seat is like billing self-checkout by the cashier.

## Why AR, Why Now?

Plot every enterprise function on a grid: how complex is the work, and how high are the stakes when it goes wrong. Complexity is the variance in inputs, the judgment required, the number of edge cases that don't pattern-match anything in training data. Stakes are what breaks when you get it wrong. A misclassified MRI is one quadrant. A poorly worded payment reminder is another.



AR sits in the bottom-left. The core question is whether someone owes us money and how long it's been outstanding. The second-order questions aren't much harder. Did the customer dispute the invoice? Did the payment match the open balance, or is there a \$0.43 gap because someone wired in the wrong currency? Polite reminder, firm reminder, or escalate? A human makes these calls in seconds, and the cost of getting it wrong is bounded. If an agent nudges the wrong customer about an overdue invoice, the worst case is mild embarrassment and a corrected record. Nobody dies and nobody gets sued.

Healthcare claims processing looks similar on paper. Read a document, match it to a record, approve or escalate. But the inputs are wildly heterogeneous (CPT codes, prior authorizations, clinical notes, plan-specific exclusions), the stakes carry regulatory weight (HIPAA, state insurance commissioners, the False Claims Act), and a wrong decision can mean a denied cancer treatment or a \$10M fraud penalty. The work is harder **and** the consequences are heavier.

This is why SOIs don't arrive in the enterprise all at once. They arrive in a sequence determined by the ratio of automation upside to error cost. Low complexity, low error cost and high language density puts AR at the front of a queue that runs through customer support, procurement, accounts payable, and most of the back

office. Functions that score worse on any axis (legal is high-stakes, R&D is high-complexity, manufacturing isn't language-mediated) wait their turn.

What gets settled in AR over the next two years sets the defaults for everything behind it: per-agent pricing tiered by role, deployment measured in days rather than quarters, agents running on top of the existing ERP, and a TAM benchmarked against payroll. But pricing AI against labor budgets is a transitional model. The anchor breaks the moment agents do work no human could (400 fraud checks in 50 milliseconds) or companies run lean enough that there's no team to benchmark against.

The companies winning in AR are proving the SOI model works at the intersection of low-complexity work, high-volume language, and structured enterprise data. Once that proof is in, the cost of capital for the next category drops, the credibility transfers, and the motion points at the next function down the list. The mid-market AR clerk reaching for Excel isn't a niche persona. She's the leading indicator.

For founders, the window is open to start defining category defaults. For investors, it's time to rethink your TAM models. And for buyers, the question worth asking now is how you define value when the work no longer maps to your employee's 9-to-5.

---

<sup>1</sup> Mordor Intelligence, [ACCOUNTS RECEIVABLE AUTOMATION MARKET SIZE & SHARE ANALYSIS - GROWTH TRENDS AND FORECAST \(2026 - 2031\)](#), 2026

<sup>2</sup> Ibid

<sup>3</sup> ZIPPIA, [Accounts receivable clerk demographics and statistics in the US](#), 2026

<sup>4</sup> Robert Half, [Accounts Receivable Clerk](#), 2026

**Disclaimer:** The information contained herein is provided for informational purposes only and should not be construed as investment advice. The opinions, views, forecasts, performance, estimates, etc. expressed herein are subject to change without notice. Certain statements contained herein reflect the subjective views and opinions of Activant. Past performance is not indicative of future results. No representation is made that any investment will or is likely to achieve its objectives. All investments involve risk and may result in loss. This newsletter does not constitute an offer to sell or a solicitation of an offer to buy any security. Activant does not provide tax or legal advice and you are encouraged to seek the advice of a tax or legal professional regarding your individual circumstances.

This content may not under any circumstances be relied upon when making a decision to invest in any fund or investment, including those managed by Activant. Certain information contained in here has been obtained from third-party sources, including from portfolio companies of funds managed by Activant. While taken from sources believed to be reliable, Activant has not independently verified such information and makes no representations about the current or enduring accuracy of the information or its appropriateness for a given situation.

Activant does not solicit or make its services available to the public. The content provided herein may include information regarding past and/or present portfolio companies or investments managed by Activant, its affiliates and/or personnel. References to specific companies are for illustrative purposes only and do not necessarily reflect Activant investments. It should not be assumed that investments made in the future will have similar characteristics. Please see "full list of investments" at [activantcapital.com/companies/](http://activantcapital.com/companies/) for a full list of investments. Any portfolio companies discussed herein should not be assumed to have been profitable. Certain information herein constitutes "forward-looking statements." All forward-looking statements represent only the intent and belief of Activant as of the date such statements were made. None of Activant or any of its affiliates (i) assumes any responsibility for the accuracy and completeness of any forward-looking statements or (ii) undertakes any obligation to disseminate any updates or revisions to any forward-looking statement contained herein to reflect any change in their expectation with regard thereto or any change in events, conditions or circumstances on which any such statement is based. Due to various risks and uncertainties, actual events or results may differ materially from those reflected or contemplated in such forward-looking statements.