o2c lab
Consulting & Advisory

Paraglide

# AGENTIC AI IN ORDER-TO-CASH:

## How To Tell Hype From Reality, and Where The Real Use Cases Are

0330 0884828 enquiries@o2clab.com www.o2clab.com

# Agentic AI In Order-to-Cash: How To Tell Hype From Reality, And Where The Real Use Cases Are

***A session delivered to the O2C Lab community by Rasmus Areskoug, Founder and CEO of Paraglide.***

AI agents are reshaping **order-to-cash.** They reply to billing queries, run two-way collections conversations, parse remittances, and manage supplier portals. They also support credit approvals and self-billing reconciliations.

This is the repetitive, communication-heavy work that credit controllers and AR analysts have always handled by hand, and it is work no previous generation of software could touch.

O2C runs on messy, unstructured inputs: free-text emails, exceptions, customers who reply with a question. That is exactly the kind of work agents are built for.

Think of agents as part of the finance team. They do the work; people approve and manage it.

AI-native solutions, such as **Paraglide**, put agents at the centre rather than bolting them onto older software systems.

## Key Takeaways

- ✓ AI agents are goal-oriented, tool-using systems that complete multi-step work, not rules-based automation or a chatbot bolted onto existing software.
- ✓ Finance teams spend 30% to 50% of their time on manual work, yet lag engineering, support, and sales in AI adoption.
- ✓ The strongest O2C use cases are the high-volume, conversational tasks: billing queries, collections follow-ups, cash application, deductions, credit gathering, and supplier portal actions.
- ✓ Safe deployment means starting in suggest mode and moving to act mode gradually, beginning with simpler tasks and smaller customers.
- ✓ **Paraglide** is a leading AI-native solution for order-to-cash and accounts receivable

o2c lab
Consulting & Advisory

Paraglide

# What Is An AI Agent? How Is It Different From Automation?

An AI agent is an independent system that completes work toward an objective, rather than following a fixed script. You give it a goal, and it reasons through the steps, uses tools to act, and works within the limits you set. Four properties define it:

- **Goal-oriented:** you give the agent an objective, not a flowchart.
- **Tool-using:** it can send emails, update systems, and take actions in other software.
- **Autonomous, within limits:** it acts on its own, but you regulate how much autonomy it has.
- **Composable:** It can be divided into sub-agents that each handle part of a process.

Agents are also capable of multi-step reasoning, which is what separates them from earlier waves of technology. To see why that matters, it helps to place agents against what came before.

| Technology | What It Does | Where It Struggles |
|---|---|---|
| Rules-based automation (RPA, no-code workflows) | "If this, then that" logic. Automates structured tasks like invoice posting and reconciliations | Becomes endlessly complex when there are many possible outcomes; requires structured data |
| Machine learning | Learns patterns from data to forecast, classify, and score without explicit rules | Predicts and scores, but does not carry out multi-step work |
| Generative AI co-pilots (LLMs) | Understands and generates text; good for summarising, explaining, drafting, and answering questions on data | Assists a human rather than completing the work itself |
| AI agents | Goal-oriented, tool-using systems that reason through multi-step work and act in real systems | Needs guardrails and oversight, especially early in deployment |

The shift agents introduce is structural. Before AI, humans did the work and software helped them track it. With agents, the agent does the work and the human moves up a layer, approving actions and monitoring communication. The software layer becomes an agentic layer.

## Why Are Finance Teams Behind On AI Agents?

Finance teams are behind because the core of accounts receivable is conversational, and conversation was the last thing automation could handle. More than 90% of software developers now use AI agents, and customer support has already been reshaped by them. Finance has not moved at the same pace, despite spending 30% to 50% of its time on manual, repetitive work.

Part of the reason is that earlier automation only reached the structured edges of finance. RPA automated data entry, but humans still handled the exceptions and the conversations. Intelligent document processing opened up new automation in accounts payable, accounting, reconciliations, and fraud monitoring, where the inputs are comparatively structured. AR is different. Resolving a billing query, chasing a payment, or negotiating a deduction means reading free-text email, understanding account history, and replying in context. That work stayed with humans because nothing else could do it.

Agentic AI changes that, because it unlocks contextual automation: automation that reads the situation and responds appropriately, rather than firing a fixed rule. As the Stanford 2026 AI Index Report puts it, the AI you are using today is the worst one you will ever use. The underlying models keep improving, which means tasks that were impossible six months ago are becoming routine.

---

## How Do You Tell AI Hype From Reality?

Start from pains and use cases, not from the technology. If you are solving a problem no one cares about, adding AI does not help. The most useful question to ask your team is simple: what are the most manual processes we have never been able to automate? Those are the candidates, because what is possible now was not possible six months ago.

The second filter is architecture. There is an enormous amount of AI-washing in the market, where existing products add a thin layer of AI to a system that was never built for it. Moving a pre-AI product to genuine agentic capability is as hard as moving from on-premise to the cloud. When that shift happened, cloud-native software held the advantage because it was built for the new model from the start. The same pattern is repeating now. AI-native products such as **Paraglide** are built entirely differently, with agents in mind from the beginning, on architecture that assumes the capabilities now available.

So when you evaluate tools, look for AI-native solutions, or build the capability yourself. A feature labelled "AI" on top of legacy automation is not the same thing as a system designed around agents. The strongest use cases in O2C are the high-volume, conversational tasks in credit control that earlier automation could never reach. Paraglide runs a team of AI agents across the order-to-cash process, each handling a distinct part of the work.

o2c lab
Consulting & Advisory

Paraglide

# Where Are The Agentic Use Cases In Order-to-Cash?

| Use-case | What AI Agents Do |
|---|---|
| **Invoice inquiry management** | Works in your finance inbox 24/7, resolving billing queries and disputes: invoice and billing questions, document requests, payment status questions, disputes and deductions, remittance parsing, PO number updates, and master data updates pushed to the ERP |
| **Collections management** | Sends personalised payment reminders and manages replies and follow-ups within existing threads using full conversation context. Escalates to the right contact at the right time, reroutes when emails bounce, tailors strategy by segment, and logs all promise-to-pay dates and commitments |
| **Credit management** | AI agents handle credit approvals and customer onboarding, combining external credit score with internal payment behaviour |
| **Cash application** | Traditional rules-based matching engines break when the transaction and invoice don't match directly, leaving edge cases to be handled manually by humans. AI agents can read and parse remittances in any format and match the payments that rule-based engines reject: wrong references, missing digits, short payments, and bulk transfers. If the customer hasn't sent a remittance, overpaid or deducted an amount without clarifying, the AI agent can reach out directly to the customer to ask for clarification. Once received, the AI agent captures the evidence and matches the payment accordingly |
| **Dispute management** | Disputes that get lost in the finance inbox, without visibility or accountability, are a frequent issue that is blocking payments. AI agents detect disputes directly from your finance inbox, gather context across systems, escalate to the right stakeholders, and follow up until resolved |
| **Supplier portal AI agent** | Maintains secure access to supplier portals and logs in through the browser to act: uploads invoices and maps them to the correct PO number, enters required metadata, tracks invoice status across portals, resolves rejections, and downloads documents |
| **Dunning** | AI agents personalise contextual reminders and continue to follow up in existing threads like a human would |
| **Deduction management** | AI agents parse remittances, capture deduction codes, and reach out to customers if deduction reasons are missing |

o2c lab
Consulting & Advisory

Paraglide

# How Do You Deploy AI Agents Safely?

Deploy agents safely through gradual automation: start in suggest mode and move to act mode over time. The progression has three stages:

1) **Humans do the work:** The starting point, where software helps out but the people carry out every action that is necessary.

2) **Agents suggest, humans approve:** The agent proposes the action or the reply, and a human approves it before anything happens. This is suggest mode.

3) **Agents act, humans monitor:** The agent takes action on its own, and humans oversee rather than approve each step. This is act mode.

The safe path is to expand autonomy along two axes at once: task complexity and customer size. Begin with simpler tasks for smaller customers, prove it works, then roll out more broadly.

| Start Automating Early (Simpler Tasks) | Keep Humans In The Loop Longer (Complex Tasks) |
|---|---|
| **Collections follow-ups** | Master data updates such as billing details |
| **Re-sending documents** | Disputes, claims, and refunds |
| **Billing clarifications** | Credit decisions |
| **Email bounces** | Contractual changes |
| **Promise-to-pay logging** | Collections escalations such as any suspension of service and credit adjustments |

The same logic applies across your customer base. For high-value customers, agents suggest replies and actions while humans review and approve the communication. For the long tail, agents handle communication automatically and only escalate critical tasks such as disputes and refunds. You decide where agents act and where humans do, based on task complexity and customer value.

This is also where flexibility matters. The strongest setups combine agentic workflows with rules-based ones and human-in-the-loop logic, rather than forcing everything through one model. In Paraglide, autonomy levels are configured directly in the agent studio, so the boundary between suggest and act can be set task by task and segment by segment. The direction of travel is clear. In 2024, humans did the work. In 2025, agents began to suggest it. In 2026, agents do it. The job of an AR user in 2026 is to manage a team of agents.

## What Does Going Live Actually Look Like?

Getting Paraglide agents into production follows a structured path that reaches go-live in around ten days and then expands autonomy from there.

**Week 1: Integration Setup and Onboarding**
Integrations are set up to pull in invoices, payment status, customers, and contact details. The agent is connected to your email inbox, workflow cadence and segments are built, and fields are mapped between systems.

**Week 2: Agent Training and Testing**
A knowledge base is built to give the agent context on how to reply to customer queries. The agent is given its tools (email, SMS, API access, browser access), custom prompts are written as instructions, and guardrails and human-in-the-loop logic are set.

**Week 3-4: User Onboarding and Go-Live**
Users are onboarded into managing agents through an approval-based workflow. The system goes live in suggest mode, with humans approving the tasks the agents propose.

**Week 3-10: Gradual Automation**
Agents improve over time and are gradually given more autonomy to act without human involvement, starting with long-tail customers and simple queries.

**Week 10 and Beyond: Autonomous Agents**
Agents act autonomously and can take action in other systems, with humans approving and gradually extending autonomy.

---

## What Results Do AI Agents Deliver In O2C?

The clearest result is faster cash. One global wholesaler reduced total AR outstanding by 34% within two weeks of going live with Paraglide.

As Tobias Pfleger, Finance Director at Choco, described it: "We saw a direct impact on our cash flow when we deployed Paraglide. Within 14 days, we had already reduced DSO by 34%. We went live quickly with minimal effort from our team."

**Across deployments, the impact shows up in four areas:**

| Metric | Result |
|---|---|
| **Reduction in DSO** | 34% |
| **Less manual work** | 42% |
| **Reduction in bad debt** | 26% |
| **Customer satisfaction** | Improved NPS |

**Those numbers come from a few connected mechanisms:**

- ✓ **Faster DSO**, because agents resolve customer queries faster, prioritise tasks by impact on DSO, send personalised messages that fewer customers ignore, identify bad contact data early, and fully automate the long tail you could never staff.
- ✓ **Less bad debt**, because agents solve payment-blocking issues sooner, catch overdue invoices as early as possible, flag bad contact data and escalate to new contacts immediately, and run automated escalation flows such as service suspension, legal letters, or referral to collection agencies.
- ✓ **Less manual work**, because agents remove the high-volume, repetitive tasks (common queries, follow-ups, reminders, payment chasing) so you can scale revenue without growing the AR team and focus people on strategic work.
- ✓ **Better relationships**, because segmented, personalised collections let you get paid on time without straining customer goodwill, turning off automatic escalations for high-value accounts and bringing account managers into the process where it counts.

---

## Frequently Asked Questions

### What is the difference between an AI agent and rules-based automation in AR?

Rules-based automation follows fixed "if this, then that" logic and needs structured data, which makes it brittle when there are many possible outcomes. An AI agent is goal-oriented and uses tools to complete multi-step work, reasoning through situations such as a free-text billing query rather than matching a predefined rule.

o2c lab
Consulting & Advisory

Paraglide

### Why have finance teams been slower than other functions to adopt AI agents?

Accounts receivable work is conversational, and conversation was the hardest thing for earlier automation to handle. While more than 90% of developers use AI agents and support has already been transformed, AR tasks like resolving queries and chasing payments stayed with humans because nothing else could read context and reply appropriately until recently.

### What does AI-native mean, and why does it matter for AR software?

AI-native means the product was built around agents from the start, on architecture that assumes current AI capabilities, rather than adding an AI feature to a legacy system. The shift from pre-AI software to agentic software is as hard as the move from on-premise to cloud, and AI-native products hold the same advantage cloud-native products once did.

### How do you deploy AI agents in accounts receivable without losing control?

Start in suggest mode, where agents propose replies and actions for humans to approve, then move to act mode gradually. Expand autonomy along task complexity and customer size together, beginning with simpler tasks for smaller customers, and keep humans in the loop for complex work such as disputes, credit decisions, and key accounts.

### Which O2C tasks should you automate first?

Begin with high-volume, lower-risk tasks: collections follow-ups, re-sending documents, billing clarifications, email bounces, promise-to-pay logging, remittance parsing, and PO number capture. Hold humans in the loop longer for master data updates, disputes and refunds, credit decisions, contractual changes, and collections escalations.

### How long does it take to get AI agents live in accounts receivable?

With Paraglide, integration and onboarding happen in week one, agent training and testing in week two, and go-live in suggest mode within weeks three to four, reaching production in around ten days. Autonomy is then expanded gradually, with agents acting more independently from week ten onward.

### What results can AI agents deliver in order-to-cash?

One global wholesaler cut total AR outstanding by 34% within two weeks of going live. Across deployments, Paraglide is associated with a 34% reduction in DSO, 42% less manual work, a 26% reduction in bad debt, and improved customer satisfaction scores.

## Summary

Agentic AI is the first technology able to handle the conversational core of accounts receivable, which is why O2C is one of its strongest applications and also why finance has been slow to adopt it. The way to separate hype from reality is to start from your most stubborn manual processes and to insist on AI-native architecture rather than AI bolted onto legacy automation.

The highest-value use cases are the high-volume, contextual tasks across billing support, collections, reconciliation, credit, and supplier portals. And the safe way to deploy is gradual: suggest mode first, act mode later, expanding from simple tasks and smaller customers outward. By 2026, the job of an AR team is less about doing the work and more about managing a team of agents that do it.

Paraglide

**Automate collections and invoice queries with AI agents**

Get paid faster with AI agents that manage collections and inbound invoice queries end to end in your finance inbox, resolving disputes and capturing key payment data like promise-to-pay dates and PO numbers.

*Paraglide builds AI agents for accounts receivable, working across the full order-to-cash process. Learn more at paraglide.ai.*

**Web Address: https://www.paraglide.ai/**

# Get In Touch

**We are building something great - and we want you to be part of it.**

**Phone: 0330 0884828**

**Email: enquiries@o2clab.com**

**Web Address: https://o2clab.com/**

**SCAN ME**