WHITE PAPER

The 2026 AI Inflection Series - Chapter 13

How Computer-Using Agents Will Transform Enterprise Execution

The next AI advantage will not come from chat alone. It will come from agents that can act inside the messy software environments where enterprise work still stalls.

> Computer use has moved from research novelty into live platform direction. OpenAI introduced Operator and later integrated that direction into ChatGPT agent. Anthropic continued advancing computer use and acquired Vercept in February 2026 to deepen the capability. This chapter examines why that matters for the real enterprise. Not the clean API story. The interface story.

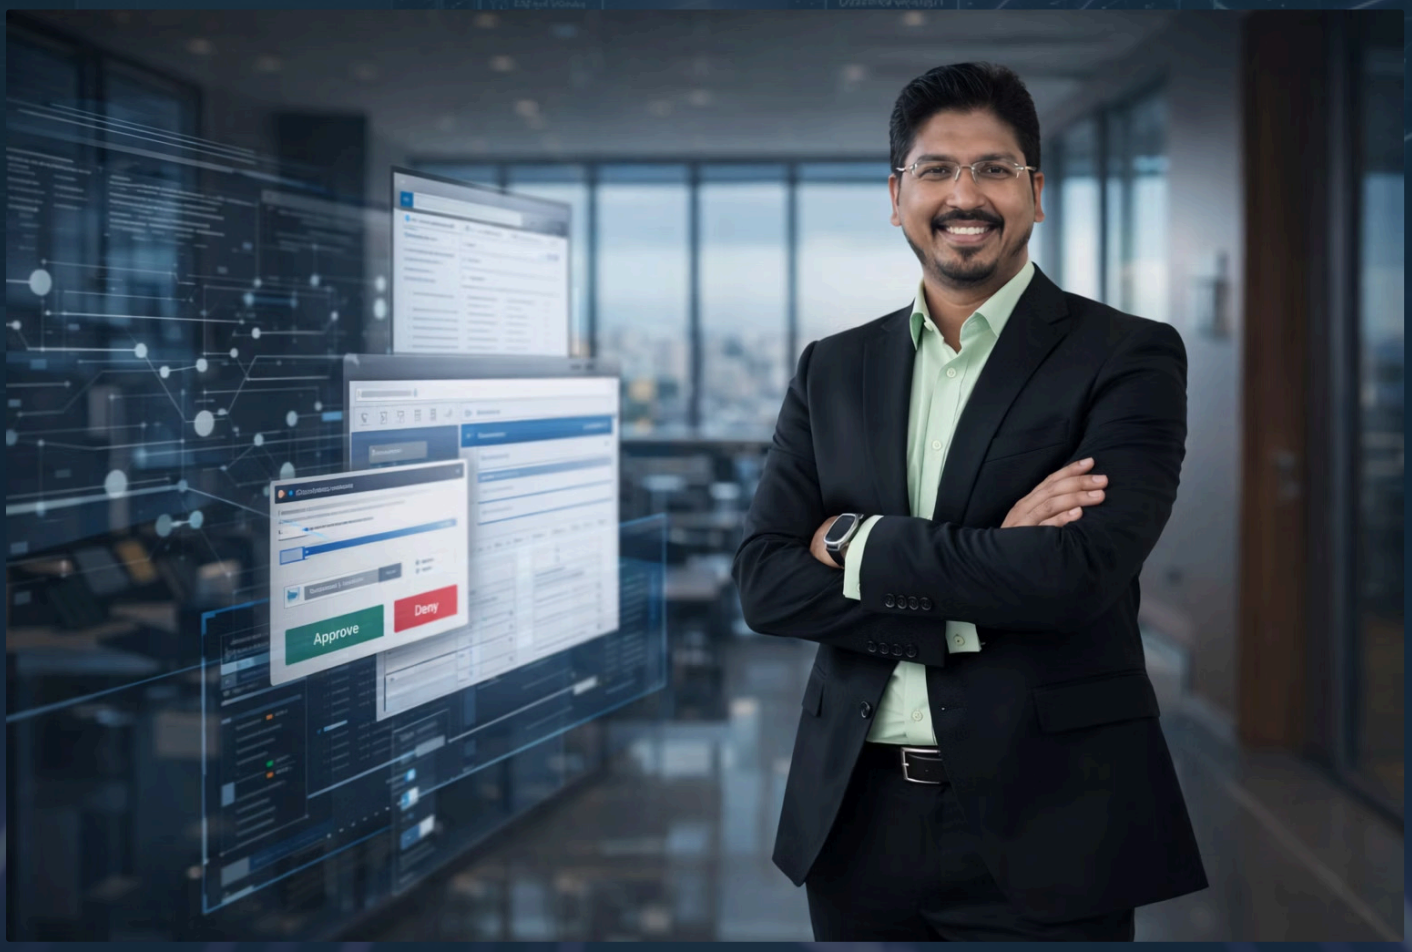

By: Logan Sivanasen | 7th April 2026

Executive Summary

This chapter examines how computer-using agents are moving from research novelty into enterprise execution and what that shift demands from leaders in governance, operations, and commercial strategy.

The Problem

Enterprise AI is stalling at the execution layer, not the model layer. Work still moves through portals, spreadsheets, approval chains, and legacy interfaces that no API or chat interface can reach. That gap is where value leaks.

The Signal

OpenAI launched Operator in January 2025 and folded it into ChatGPT agent by July. Anthropic pushed OSWorld scores from under 15% to 72.5% and acquired Vercept in February 2026. The benchmark-to-product cycle is compressing fast - this is now a platform race.

The Opportunity

Computer-using agents operate inside live interfaces the way a person at a keyboard would - without requiring clean APIs. The first wins will come from narrow, high-frequency, rules-bounded workflows: invoice matching, CRM hygiene, procurement portals, claims intake. Supervised execution, not open-ended autonomy.

The Risk

Prompt injection is the most significant hidden risk. Once an agent holds live credentials and browses real systems, adversarial content can hijack execution. The control stack - policy, permissions, confirmation gates, logging, containment - is not overhead. It is the foundation.

The Action

Start narrow. Instrument from day one. Scale with evidence, not enthusiasm. Map last-mile workflows, score by friction and risk, pilot two or three with mandatory human checkpoints, and build the governance architecture before expanding. The window to build institutional capability is open - but not indefinitely.

The shift is not from human work to machine work. The shift is from stalled execution to supervised, instrumented execution.

Sources: OpenAI API Docs, "Computer use." Anthropic, "Anthropic acquires Vercept," 2026. OpenAI, "Introducing ChatGPT agent," 2025. NIST AI RMF resources.

This category crossed from demo to direction

The shift is now visible in both product and benchmark signals. What began as a research preview is now a declared platform direction across the two most influential AI labs in the world.

Jan 2025

OpenAI launches **Operator**. The model can go to the web and perform tasks for users.

Mar 2025

OpenAI releases **developer computer-use tools**. Reports 38.1% on OSWorld for full computer-use tasks.

Jul 2025

OpenAI folds Operator into **ChatGPT agent**. Computer use becomes core product direction.

Feb 2026

Anthropic reports **72.5% on OSWorld**, up from under 15% in late 2024. Acquires Vercept to deepen computer-use capability.

38.1%

OpenAI OSWorld

Full computer-use task score, developer tooling release, March 2025

72.5%

Anthropic OSWorld

Score by February 2026, up from under 15% in late 2024. Described as approaching human-level performance.

14

Months

From Operator launch to Vercept acquisition. The benchmark-to-product cycle is compressing fast.

Research novelty is over. Execution capability is now a platform race.

Sources: OpenAI, "Introducing Operator," 2025. OpenAI, "New tools for building agents," 2025. Anthropic, "Anthropic acquires Vercept," 2026.

Transformation still breaks inside interfaces

The promise of AI often gets framed at the model layer. The real drag sits lower. It sits in quote approvals, CRM hygiene, procurement portals, claims systems, finance reconciliations, partner platforms, service consoles, and spreadsheet-based coordination. These are not edge cases. They are the daily execution layer of the enterprise.

When APIs are incomplete or absent, work defaults back to human copy-paste, context switching, and delay. No model-layer improvement resolves this. No conversational interface bridges it. Computer-using agents matter because they can operate in that exact layer.

> This is the bridge between AI promise and the ugly reality of how work still gets done. Computer-using agents do not require clean APIs. They work through the interface, the way a person does.

Operational Friction

- Quote approvals and finance reconciliations stall at the interface layer
- Partner platforms and service consoles remain disconnected from AI tooling

Manual Workflows

- Without clean APIs, work defaults to human copy-paste and context switching
- Spreadsheet-based coordination persists as the daily execution reality

Broken Integrations

- Procurement portals and legacy systems have no API surface for AI to reach
- Any model-layer improvement stops at the interface boundary

Human Copy-Paste

- Manual data entry and re-keying across systems consumes execution capacity
- Context is lost at every handoff between disconnected tools

AI Promise

Model-layer improvements announced

Interface Reality

Work still routes through portals, spreadsheets, approvals

Execution Gap

No API, no bridge - humans fill the gap manually

Agent Fix

Computer-using agents operate at the interface layer directly

Sources: OpenAI API Docs, "Computer use." Anthropic, "Anthropic acquires Vercept," 2026.

Work is overloaded. Customers moved faster than systems did.

Microsoft's 2025 Work Trend Index surfaced a capacity gap that leaders can no longer ignore. The numbers are not marginal. They describe a structural mismatch between what organizations expect to produce and what the current operating model can actually sustain.

53%

Productivity mandate

of leaders say productivity must increase, with no corresponding increase in capacity

80%

Capacity shortage

of the global workforce says they lack enough time or energy to do their work

82%

Digital labor intent

of leaders expect to use digital labor to expand workforce capacity in the next 12 to 18 months

81%

Agent integration

of leaders expect agents to be moderately or extensively integrated into their AI strategy over the same window

The workday itself is fragmenting at the attention layer. Microsoft's "infinite workday" reporting shows employees using Microsoft 365 are interrupted every two minutes by meetings, email, or notifications. The average employee receives 117 emails and 153 Teams messages daily. That is not a productivity problem. That is an execution architecture problem.

117 emails

Average daily email volume per employee in Microsoft 365 environments

153 messages

Average daily Teams messages per employee, compounding context-switching costs

Every 2 min

Average interruption frequency for employees, fragmenting deep work and execution continuity

83% of marketers

say customers now expect two-way conversations, while 69% struggle to respond with the context they need (Salesforce)

Microsoft Work Trend Index, 2025. Microsoft, "Breaking down the infinite workday," 2025. Salesforce, "State of Marketing 2026," 2026.

The first wins will come from narrow, repetitive, high-friction workflows

The strongest early use cases will not be open-ended autonomy. They will be supervised execution in repetitive, rules-bounded workflows with visible exceptions. OpenAI explicitly positions computer use for browser-based workflows such as QA and data entry across legacy systems. Anthropic describes computer use as enabling multi-step tasks in live applications and solving problems impossible with code alone.

Think quote validation, invoice matching, CRM updates, form completion, service routing, claims intake, campaign operations, spreadsheet refreshes, evidence gathering, and portal maintenance. The pattern across all of them is the same: high frequency, clear rules, measurable exceptions, and a human checkpoint that matters.

Finance Operations

Invoice matching, reconciliation, and approval routing inside legacy finance portals

CRM Hygiene

Contact updates, opportunity staging, and data enrichment across CRM interfaces without API dependence

Procurement Portals

Vendor form submission, purchase order entry, and status checking across supplier portals

Claims and Service

Claims intake, eligibility checking, routing logic, and service console updates at volume

Best first pilots share these characteristics

1. **High frequency.** Volume justifies the control investment.
2. **Low ambiguity.** Rules are clear enough to define exceptions explicitly.
3. **Clear business value.** Cycle time, throughput, or error rate is measurable.
4. **Limited permission scope.** The agent does not need broad system access.
5. **Measurable exception rate.** You can instrument when the agent needs help.
6. **Mandatory human checkpoints.** Defined escalation paths, not optional guardrails.

Sources: OpenAI, "New tools for building agents," 2025. Anthropic, "Anthropic acquires Vercept," 2026.

Computer-using agents add judgment to interface action

Traditional automation worked best when steps stayed stable, deterministic, and tightly predefined. A changed field label, a new screen layout, or an unexpected modal could break an entire automation path, and that brittleness has constrained RPA adoption at scale for years.

- Field labels, screen layouts, and modal dialogs change without notice
- Exception handling requires human rewriting of the entire automation path
- No ability to recover context when an unexpected state is encountered
- Tightly scoped to one interface version; any update can cause silent failures

Computer-using agents push further because they can interpret live screen state, navigate changing interfaces, and act across multi-step tasks with contextual flexibility. OpenAI describes computer use as the model inspecting screenshots and returning interface actions - reporting leading scores on OSWorld, WebArena, and WebVoyager. That does not make the category fully reliable. It does make it materially more capable than brittle click-path automation alone.

- Interpret dynamic UI elements and recover from unexpected screen states
- Navigate interface variation without requiring exact screen state matches
- Flag ambiguous edge cases for human review rather than failing silently
- Escalate with full context, not just error codes

Legacy Automation

- Brittle click-path sequences
- Breaks on interface change
- No contextual recovery
- Limited exception handling
- Requires exact screen state
- Human rewrites automation on failure

Computer-Using Agent Execution

- Inspects live screen state via screenshot
- Navigates interface variation adaptively
- Contextual recovery on unexpected states
- Flags edge cases for human review
- Interprets dynamic UI elements
- Escalates with context, not just error codes

> The shift is not from automation to autonomy. The shift is from brittle automation to supervised adaptive execution.

Sources: OpenAI API Docs, "Computer use." OpenAI, "New tools for building agents," 2025.

The control layer becomes the real product

Once an agent can act inside live systems, governance stops being a policy appendix. It becomes part of execution itself. NIST's AI Risk Management Framework and Generative AI Profile give organizations a governance structure for identifying and managing AI-specific risks. OpenAI's agent deployment and safety materials make the practical version clear: use explicit confirmation for consequential actions, use active supervision for sensitive tasks, narrow the scope of access, limit what data the model can see, log behavior, and review failures.

This is not a compliance exercise. It is an operating architecture decision. Organizations that treat the control stack as an afterthought will discover it is actually the foundation. The control layer is what makes supervised execution trustworthy enough to scale.

1

Policy

What the agent is allowed to do. Defined boundaries. Explicit task scope. No open-ended access grants.

2

Permissions

What the agent can access. Scoped credentials. Allowlists. Session boundaries. Data visibility limits.

3

Execution

Which actions require human approval before the agent proceeds. Confirmation logic for consequential steps.

4

Observation

What gets logged and reviewed. Audit trail. Anomaly detection. Failure review cadence.

5

Containment

What stops a bad path from spreading. Rollback capability. Environment separation.

The control stack is not overhead. It is what transforms a capable agent into a deployable one. Every layer compounds trust.

Sources: NIST AI RMF and GenAI Profile. OpenAI, "Introducing ChatGPT agent," 2025.

The biggest hidden risk is unintended action

As soon as an agent browses, reads documents, or works through live systems, prompt injection becomes more than a model-quality issue. It becomes an execution risk. OpenAI says prompt injection is one of the most significant risks for browser agents and describes it as a long-term challenge for agent security. Anthropic says computer use also poses risk because malicious actors can hide instructions on websites to hijack the model. Both companies now treat this as a core safety and deployment issue, not a side note.

That means leaders must stop asking only whether the model is smart enough. The harder question is whether the control stack is strong enough if the agent sees adversarial content while holding valid access to real enterprise systems. An agent that is operating inside a CRM, a finance portal, or a procurement system with live credentials is not just a productivity tool. It is an execution actor with access that can be weaponized if the control architecture has gaps.

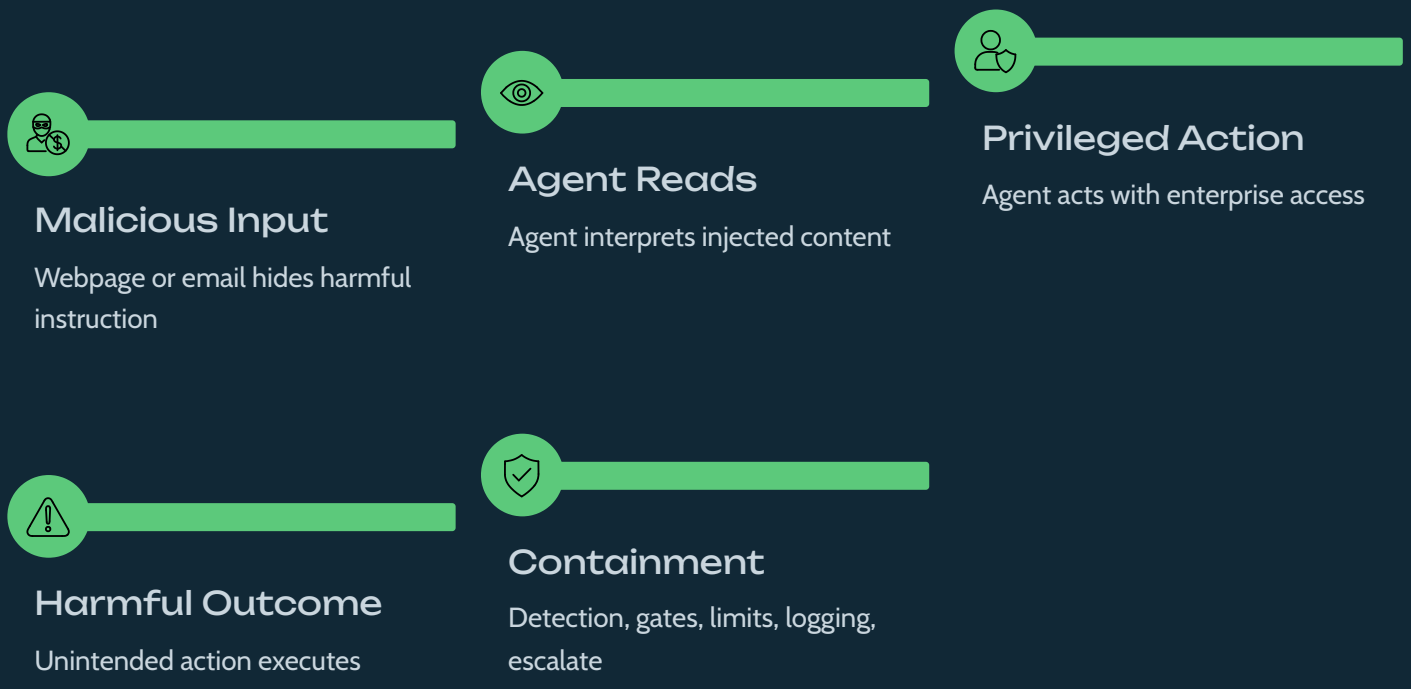

Detection. Scan content before the agent acts. Flag anomalous instruction patterns.

Confirmation. Require human approval before any action triggered by external content.

Restricted permissions. The agent should never hold more access than the specific task requires.

Logging. Every action must be attributable, timestamped, and reviewable post-execution.

Escalation. Define what triggers a stop-work event versus a review-required flag.

Sources: OpenAI, "Hardening ChatGPT Atlas against prompt injection," 2025. Anthropic, "Anthropic acquires Vercept," February 25, 2026.

Execution speed now shapes revenue

This chapter is not only about productivity. It is about commercial performance. IBM's 2025 CMO research surfaces alignment gaps that translate directly into revenue loss. When execution is slow, fragmented, or disconnected from customer signals, the commercial cost is not theoretical.

28%

CX ownership alignment

of surveyed organizations say the end-to-end customer experience is effectively owned and aligned across functions (IBM, 2025)

22%

AI decision guardrails

have clear guidelines and guardrails for the use of AI in automated decision-making (IBM, 2025)

20%

Revenue unlock

potential increase in revenue from fully aligning marketing, sales, and operations (IBM CMO study, 2025)

When customer expectations move faster than the operating model, revenue leaks through execution delay. The lead that did not get a timely response. The quote that sat in an approval queue. The service case that lost context across handoffs. Computer-using agents matter because they target the places where demand, response, fulfillment, and internal action still break apart.

Before: Manual Execution Flow

1. Inbound lead or customer signal arrives
2. Manual routing to the right team or system
3. Delay from queue depth and context switching
4. Missing context across disconnected systems
5. Lost momentum and degraded customer experience

After: Agent-Assisted Execution

1. Signal received and classified by agent in real time
2. Agent executes bounded steps inside live systems
3. Human checkpoint fires on high-impact actions
4. Faster response with full context preserved
5. Cleaner handoff and measurable commercial outcome

Sources: IBM, "Profit-Driven CMOs See AI as Growth Driver, but Operational Hurdles Slow Them Down," 2025.

The enterprise will need agent supervision, not only agent access

Computer-using agents do not remove the need for people. They raise the standard for workflow ownership, escalation design, and operational judgment. Microsoft's Work Trend Index says 78% of leaders are considering hiring for AI-specific roles, with roles under consideration including AI trainers, security specialists, AI agent specialists, ROI analysts, and AI strategists. The same report says leaders are already thinking in terms of digital labor and hybrid teams of humans and agents working together.

In practice, the new jobs around computer use will center on workflow ownership, permission design, exception review, evals, and policy enforcement. Human work moves upward toward judgment, oversight, redesign, and accountability. This is not workforce reduction. It is workforce reorientation. The people who thrive will be the ones who can design, govern, and continuously improve the workflows agents execute.

Old Operating Model

Employee owns and completes the full manual execution path. Context held in their head. Exceptions resolved informally. No instrumentation. Throughput is a function of headcount.

New Operating Model

Agent executes bounded steps inside live systems. Human reviews exceptions and edge cases. Workflow owner designs, monitors, and improves the execution architecture. Throughput is a function of workflow design quality.

Workflow Ownership

Someone must own the end-to-end design, exception logic, and continuous improvement of every agent-assisted workflow

Permission Design

Scoping agent access is a skilled operational task, not a one-time IT checkbox. It evolves with the workflow

Exception Review

Human reviewers become accountable for the cases agents escalate. That is higher-stakes work, not lower

Evals and Impact

Measuring throughput, intervention rate, and commercial outcome is the new operational reporting layer

Sources: Microsoft Work Trend Index, 2025.

Start narrow. Instrument hard. Scale with evidence.

The right rollout path is practical and sequenced. The organizations that will get ahead are not the ones moving fastest. They are the ones moving most deliberately, with clear instrumentation and defined escalation from the first day of the pilot. This sequence aligns with the control logic in current OpenAI safety guidance and NIST risk-management framing.

1

Map

Document last-mile workflows. Rank by friction, business value, exception risk, and system sensitivity. Identify where execution actually breaks, not where leaders think it breaks.

2

Score

Prioritize using four dimensions: task frequency, permission exposure, exception ambiguity, and recoverable failure cost. Disqualify high-ambiguity workflows from the first cohort.

3

Pilot

Select two or three workflows. Define mandatory human checkpoints before you deploy a single agent step. Build the control structure first, then the capability.

4

Control

Review prompt-injection exposure before scale. Define what constitutes a recoverable failure, a review-required failure, and a stop-work event. Log everything from day one.

5

Measure

Instrument throughput, completion rate, intervention rate, exception rate, recovery time, and business impact. Make the evidence visible before any expansion conversation.

6

Expand

Scale with evidence, not executive enthusiasm. Use pilot data to calibrate the control model for the next workflow tier. Each cohort improves the operating architecture.

Sources: NIST AI RMF and GenAI Profile. OpenAI, "Introducing ChatGPT agent," 2025. OpenAI, "Hardening ChatGPT Atlas against prompt injection," 2025.

The operator toolkit for Chapter 13

Beyond the insights in this chapter, I have created these free tools to help readers turn ideas into action. They are built to be practical, reusable, and worth keeping close as teams work through the governance, execution, and commercial challenges covered here.

Computer-Using Agent Readiness Scorecard

Score workflows by value, risk, permissions, and UI stability before committing to deployment

Last-Mile Execution Heatmap

Find where execution breaks across portals, tabs, spreadsheets, approvals, and handoffs

Human Checkpoint Design Card

Define which actions require review, by whom, and under what conditions

Prompt-Injection Control Matrix

Map exposure points, likely impacts, and containment controls for every agent-touched interface

Agent Permission and Session Blueprint

Design scoped access, allowlists, environment separation, and human-takeover rules

Computer-Use Pilot Selection Canvas

Select the best first three use cases and explicitly reject the wrong ones before resources are committed

Exception Taxonomy and Escalation Playbook

Classify recoverable failures, review-required failures, and stop-work events with clear response paths

Agent ROI Tracker

Measure throughput, completion rate, intervention rate, cycle time, and commercial impact in one view

Workflow Evidence Pack Template

Give legal, risk, and audit teams a structured view of how the workflow is governed and logged

90-Day Rollout Plan

Move from workflow mapping to pilot, evals, and scale with a structured timeline and decision gates

This page is a synthesized tool suite derived from the governance, execution, and commercial issues documented across OpenAI, Anthropic, Microsoft, IBM, Salesforce, and NIST.

The interface is the new automation frontier

The next wave of enterprise AI will not win because the model sounds smarter. It will win because work finally moves faster through the interfaces where execution still gets trapped. Computer-using agents matter because they meet the enterprise where the enterprise still is. Inside portals. Inside spreadsheets. Inside approvals. Inside fragmented systems. Not after transformation. During it.

OpenAI moved this category from research preview to product direction in under 14 months. Anthropic pushed the benchmark from under 15% to 72.5% in roughly the same window and acquired Vercept to deepen the capability. Microsoft, IBM, and Salesforce all point to the same operating tension: work is overloaded, teams are fragmented, and customers expect faster, two-way response than current systems can support. That convergence is not coincidental. It is the signal that the execution layer is where the next enterprise edge gets built.

The leaders who will be ahead of this in 24 months are not the ones waiting for the technology to mature further. They are the ones building the control stack now, running supervised pilots now, and developing the workflow ownership skills that agent-assisted execution will require at scale. The window to build that institutional capability is open. It will not stay that way.

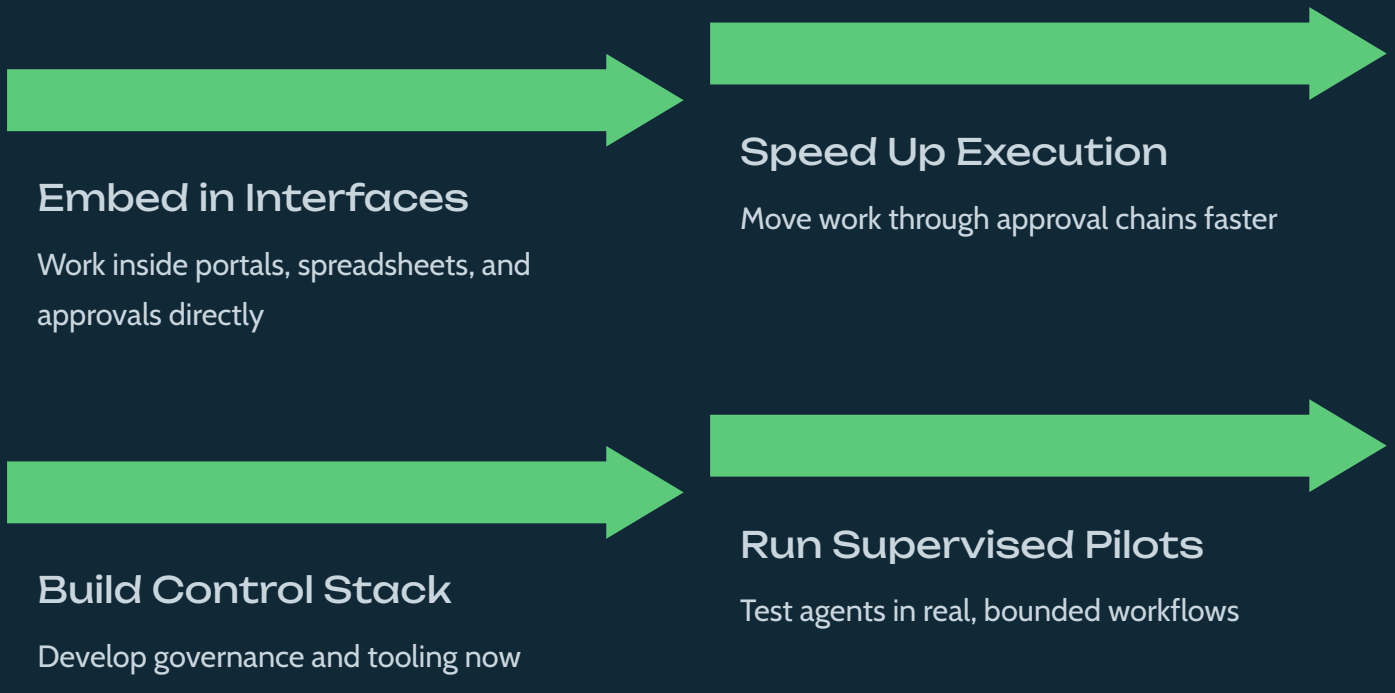

> The shift is not from human work to machine work. The shift is from stalled execution to supervised, instrumented execution.

Sources: OpenAI API Docs, "Computer use." OpenAI, "Introducing ChatGPT agent," 2025. Anthropic, "Anthropic acquires Vercept," 2026.

Sources

1

OpenAI

Introducing Operator. January 23, 2025, updated July 17, 2025.

Introducing ChatGPT agent: bridging research and action. July 17, 2025.

New tools for building agents. March 11, 2025.

Computer use. OpenAI API Docs.

Continuously hardening ChatGPT Atlas against prompt injection attacks. December 22, 2025.

2

Anthropic

Anthropic acquires Vercept to advance Claude's computer use capabilities. February 25, 2026.

Introducing Claude Sonnet 4.6. February 17, 2026.

3

Microsoft

2025: The year the Frontier Firm is born. April 23, 2025.

Breaking down the infinite workday. June 17, 2025.

New Microsoft Study Reveals the Rise of the Infinite Workday. Microsoft News. June 17, 2025.

4

IBM

Profit-Driven CMOs See AI as Growth Driver, but Operational Hurdles Slow Them Down. June 17, 2025.

5

Salesforce

State of Marketing 2026. February 19, 2026.

6

NIST

AI Risk Management Framework and Generative AI Profile resources.

This chapter is part of the 2026 AI Inflection Series, a premium executive research program examining the operational, commercial, and governance dimensions of enterprise AI at inflection. All claims are cited to primary source documents. Benchmarks reflect published figures at time of writing.