Digital Employees: The Future of AI Agent Deployment in the Enterprise

Thesis Statement
Enterprise workflows are increasingly being reimagined through experimentation with AI agents. We believe AI agents will become ubiquitous in the workplace, **empowering top performers to delegate tasks to multiple specialized digital employees working in parallel.** This will enable real-time ideation, seamless execution, and frictionless communication at a superhuman pace. Software engineers are already leading this shift—deploying AI agents that write, test, and manage code—effectively transforming a single engineer into the supervisor of multiple digital collaborators. For this paradigm to scale beyond the engineering teams and across the enterprise, organizations will need robust agent operations platforms that can create, test, deploy, orchestrate, and monitor these digital workers with the reliability, security, and compliance demanded by enterprise environments.

The right to exist for any employee not using AI agents will be questioned.

Our Definition of an AI Agent:
An AI agent is a digital employee, implemented via code, workflows, or scripts, that **operates autonomously for extended periods**, with occasional human oversight. It can access digital information and interact via human or software interfaces and **can reason** and make decisions through its connection to one or more large language models (LLMs).

Introduction
AI agents are in the early days of demonstrating that they are powerful tools for enterprise productivity. Unlike traditional bots or scripts, ***agents can reason***, use natural language, and integrate with various apps to carry out complex workflows. For example, an agent could draft **emails, update CRM entries, research information, and schedule meetings,** all based on high-level goals provided by a human. This promises to multiply employee productivity: a single employee might manage a team of AI “co-workers”, each specialized in a domain (marketing research, code generation, customer support, etc.), working in parallel to accomplish objectives.

Today, interest in enterprise AI agents is surging. Companies from startups to large corporations are experimenting with agents in roles like recruiting, customer service, and data analysis. At the same time most enterprises are proceeding with caution. Agents are still seen largely as “scripts or APIs with LLMs attached” for handling unstructured data, but not yet **trusted for mission-critical operations**. Many pilot projects are stuck in the lab, as firms grapple with agents’ unpredictable behavior (hallucinations, errors) and lack of accountability.

Bottom-up vs. top-down adoption:

There are two adoption paths in enterprises. *Bottom-up*, tech-savvy employees and teams are adopting no-code agent tools on their own to deploy digital co-workers to automate daily tasks, similar to how Slack or Zoom spread virally. This grassroots adoption is fueled by easy-to-use

agent builders, browser workflow automations, and single-use agents, where anyone with a browser and an API key can spin up an AI helper.

Top-down, CIOs and IT leaders are evaluating enterprise-grade agent platforms to deploy organization-wide and to embed AI into business processes. Enterprise leaders treating agentic AI as a strategic transformation, referring to this moment in times as the “Fifth Industrial Revolution”. Bottom-up agents drive initial experimentation and showcase quick wins, while top-down initiatives will be critical for scaling successful agent use cases firmwide (with proper governance and integration). We expect these paths to converge: pilots proven by employees will inform enterprise-wide deployments, and enterprise platforms will gradually empower end-users to customize their own agents within approved guardrails.

Status Quo

Bottom-Up Developer-Led Adoption: Early Digital Employees

Developers have become the early adopters and evangelists of AI agents co-workers inside organizations. Rather than wait for official IT initiatives, engineers and power users are **deploying “digital employees” in their own workflows**, often using whatever tools are at hand. They treat large language models as worker bots that can be instructed via natural language, essentially a **mini AI workforce** that can write code, draft content, answer questions, or execute repetitive tasks on command.

Developers leverage a variety of modalities to create and manage these personal agents: from prompt engineering, vibe coding, and scripting, to using emerging tools like AI coding copilots and automation utilities. For example, engineers now **integrate natural language tools directly into their development workflow**, using AI pair-programming editors like *Cursor* and invoking AI actions with simple shortcuts. This approach to AI “co-pilots” allows developers to automate a surprising range of work without needing permission or full-platform solutions effectively bottom-up deployment of digital coworkers.

The status quo of developer-led agent adoption is a bellwether for broader enterprise use. Today it’s the programmers chaining together prompts and scripts; ***tomorrow it will be non-technical employees*** using refined versions of these agents through user-friendly interfaces. Knowledge workers in functions like marketing, finance, and customer support are beginning to use AI assistants for drafting communications, researching information, and handling routine decisions. The modalities pioneered by developers – instructing AI in natural language, having it observe and learn from user actions, integrating it with everyday software are laying the groundwork for mass-market usage by non-technical staff. In effect, developers are the **innovation sandbox** for digital employee concepts, proving them out in code before they get packaged for everyone else.

Bottom-Up Non-Technical Adoption: Single Purpose AI Agents

We have already seen the adoption of single purpose agents by non-technical users; however, we are waiting on wide scale of adoption of orchestrating multiple single purpose agents to carry out a task autonomously like a human worker would do. Examples of single purpose agents are **Fireflies** (transcriber and summarizer), **Writer** and **Jasper** (produces copy), **Rogo** and **Gamma** (presentation creation), **Altitude** (accounting), etc. Single purpose agents solve one piece of the digital worker puzzle, the next frontier is getting them to work together activated by a high level prompt from a human.

Early adopters often accept clunky workflows or trust the AI without oversight; average employees will demand more polished, safe, and easy-to-use agents. For instance, a developer might write a custom script to have GPT analyze financial data, but a finance analyst will need that capability delivered through a secure, approved tool with guardrails. The grassroots enthusiasm from builders is there, but bridging the gap to **enterprise-wide rollouts** will require robust platforms (security, compliance, UI) that abstract away the complexity. Still, the key takeaway is that bottom-up momentum is real: *developers are already living in the future* with AI agents as digital teammates, and they're pulling the rest of the enterprise in that direction.

Today, enterprise AI agent deployment is in a nascent state. Many organizations are running proof-of-concept projects or limited pilots. Common uses include autonomous customer support bots, AI sales assistants, or internal workflow automations. These agents often start as simple scripts enhanced with LLMs – for example, a script that checks a support ticket system and uses an LLM to draft responses. This aligns with the market perception that "agents are scripts/APIs with LLMs for unstructured data processing". While these prototypes show what's possible, they also reveal current limitations.

Inflections / Why Now (please be as specific as possible, using data)

Several converging inflections make us confident that enterprise AI agents are nearing a tipping point:

- **Dramatically improved AI capabilities:** The leap from GPT-3 to GPT-4 and the emergence of other advanced LLMs in 2023–2024 have given agents much better "brains." They can handle larger context windows, complex instructions, and integrate knowledge better than a year ago. This boosts their usefulness for business tasks (e.g. reading lengthy policy documents to answer questions). Larger model sizes and fine-tuning have also made agents more capable and factual, reducing (though not eliminating) hallucinations. As one metric, OpenAI's GPT-4 scored 40% higher on factuality benchmarks than GPT-3.5 – a significant quality gain for enterprise use. The upshot: agents built on the latest models can be trusted with more complex workflows than was possible even recently.

- **Rise and Adoption of Single Purpose AI Agents:** there is steady adoption of single purpose AI Agents for non-technical work. As bottoms up adoption rise, there will be an opportunity to orchestrate several single purpose AI Agents to work together to work for extended periods of time on human prompted projects.

- **Enterprise readiness of AI providers:** The major AI platforms (OpenAI, Anthropic, Google, Microsoft) have all launched enterprise-grade offerings. With features like data privacy guarantees, audit logs, and fine-tunable models, they address CIO concerns around security and compliance. For example, OpenAI's "function calling" feature (released mid-2023) lets agents execute structured API calls more safely, reducing error rates in tool integration. Google's Duet AI and Microsoft 365 Copilot are bringing vetted AI assistance directly into Office apps. This mainstreams the concept of AI helpers for millions of business users, **conditioning them to trust AI with work tasks**. The comfort level with AI is rising at the C-suite as well, with 80% of companies expected to integrate AI agents into core operations by 2028.

- **Mounting pressure for productivity:** The current economic climate and competitive landscape push enterprises to do more with less. AI agents present a way to automate white-collar work and amplify output without proportional headcount increases. Executives see the success of early AI deployments (like coding assistants saving 30% of developers' time) and want similar leverage across all functions. This urgency has translated to budget: a recent survey showed ~55% of enterprises have allocated funds for AI automation initiatives in 2024, up from <20% two years prior (source: Gartner).

- **Emergence of the Agent Operations stack:** A wave of startups and tools is arriving to fill the gaps that made agents unreliable. We now see the outline of an "AgentOps" tooling stack analogous to DevOps for software. This includes: testing frameworks for agents (e.g. Halluminate's model-driven evaluations to catch errors), monitoring dashboards (e.g. portfolio company AgentOps, which visualizes agent steps and flags rogue behavior), simulation environments to safely train agents (e.g. Veris AI's high-fidelity "AI gyms"), and standardized libraries for connecting agents to software (e.g. hundreds of API connectors via platforms like Pipedream, Nango.dev, and Composability tools). These didn't exist a year or two ago in mature form. The lack of predictability in agent behavior is starting to be addressed by these infrastructure pieces.

- **Top-down support meets bottom-up innovation:** As described, both grassroots and executive-driven momentum is peaking now. Employees armed with new no-code agent creators are showcasing quick wins, like automating a report that used to take days. At

the same time, CEOs are tasking their tech leaders to formulate an “AI agents strategy” so the company isn’t left behind. When bottom-up innovation aligns with top-down mandate, adoption can accelerate rapidly (we saw this with cloud and with collaboration tools in the past). We expect over the next 12-18 months, many pilot agent projects will graduate to full deployments, backed by budget and oversight.

- **Data points of traction:** Lastly, we already see indicators that enterprise agents can deliver value. ***You.com***, which offers an AI research and workflow assistant, reached a $20–30M revenue run-rate and a Series C funding by 2024, demonstrating demand for AI-driven knowledge work. **LlamaIndex** (open-source project turned company) amassed 15K+ GitHub stars in months and is now used by teams at RBC, Uber, Instabase, etc., to prototype AI features over their data. Even before robust marketing, the **AgentOps** platform by Agency AI attracted “thousands of teams monthly” to monitor agents. These signals all point to a market gaining early adoption.

The demand, and the supporting ecosystem have coalesced. Enterprises have a clear incentive to deploy agents, and at the same time, they we are seeing the tools to do so in a reliable, scalable way.

Competitive Landscape

There are several types of companies that will enable AI enterprise adoption. Some of those types of companies include workflow builders, AI agent builders, monitoring and testing infrastructure, browser workflows, single purpose agents, etc. However, for non-technical users to manage a workforce of Digital Employees like engineers are doing today; there needs to be a class of companies that we refer to as All In One Agent Platforms. These platforms will have a conversational interface, allowing users non-technical users to have a vibe coding type experience when orchestrating agents to to work. There will also be an interface to see what the agents are building and achieving in real time.

These companies will ultimately become super agents, those that will orchestrate and power other agents to carry out a set of tasks originally instructed by humans.

Company	Founded	Description
Kortix	2024	Developer of a platform designed to build, host, monitor, and train artificial intelligence employees.
Athena Intelligence	2022	Developer of an AI software platform designed for artificial employees to accelerate enterprise analytics.

(645 small check)		
Cobot	2024	The first suite of AI agents built to accelerate your team.
Manus	2023	Manus is a general AI agent that bridges thoughts and actions by excelling in various tasks such as research, data analysis, education, and productivity, offering personalized solutions across different domains to help users achieve their goals efficiently.
Genspark	2023	The company's platform uses multiple specialized artificial intelligence agents to generate new, customized pages for every query, enabling users to get useful results on one page.
Abacus AI	2019	One AI assistant with access to all top LLMs, video and image generators. General purpose and coding agents included.

Market Map

You can find the ecosystem market map below. These are the companies that will enable the continued growth and adoption of AI Agents in the enterprise.

Super Agents: General Agent Assistants & All-in-One Agent Platforms

Company	Founded	Description
Kortix	2024	Developer of a platform designed to build, host, monitor, and train artificial intelligence employees.
Athena Intelligence *(645 small check)*	2022	Developer of an AI software platform designed for artificial employees to accelerate enterprise analytics.
Cobot	2024	The first suite of AI agents built to accelerate your team.
Manus	2023	Manus is a general AI agent that bridges thoughts and actions by excelling in various tasks such as research, data analysis, education, and productivity, offering personalized solutions across different domains to help users achieve their goals efficiently.
Genspark	2023	The company's platform uses multiple specialized artificial intelligence agents to generate new, customized pages for every query, enabling users to get useful results on one page.
Abacus AI	2019	One AI assistant with access to all top LLMs, video and

		image generators. General purpose and coding agents included.
ChatGPT	2022	A conversational AI interface built on OpenAI's large language models (LLMs). They launched ChatGPT Operator in 2025.

Workflow & AI Builders: code first or drag and drop

Company	Founded	Description
n8n	2019	Open-source workflow automation platform enabling users to connect apps and automate tasks via a low-code visual interface.
Ageno	2023	No-code platform for building and deploying AI-powered digital employees to automate everyday business workflows.
Langflow	2023	Visual drag-and-drop builder for designing and deploying LLM workflows on top of LangChain.
Llama Index	2022	Framework that connects LLMs to external data sources for context-aware question answering and agent development.
Crew AI	2023	Open-source framework for orchestrating collaborative, role-based multi-agent systems powered by LLMs.
Stack AI	2023	Platform that turns structured and unstructured data into custom AI applications using no-code tools and built-in LLMs.
Lindy AI	2022	AI executive assistant that automates tasks like scheduling, email drafting, and CRM updates for professionals.
Dust	2022	Developer platform for creating, testing, and deploying LLM-powered workflows and assistants for enterprise use.
Relevance AI	2020	Vector search and AI workflow platform enabling businesses to build intelligent applications using embeddings and LLMs.

AI Agent Frameworks

Company	Founded	Description

LangChain/LangGraph	2022 2023	LangChain is a developer framework for building LLM-powered applications; LangGraph is its extension for multi-agent workflows using stateful, event-driven graphs.
AutoGen (Microsoft)	2023	Microsoft's open-source framework for creating multi-agent systems where agents collaborate via conversations to accomplish complex tasks.
PromptLayer	2022	Logging and observability platform for tracking, debugging, and managing LLM prompts in production environments.
Ango AI	2021	Human-in-the-loop data labeling and annotation platform for AI model training, specializing in high-quality structured datasets.

Browser workflows

Company	Founded	Description
Automat	2023	AI-driven RPA platform that turns video demos or text descriptions into executable automations using LLMs and computer vision.
Browser use	2023	Lightweight open-source project that enables agents to operate web browsers to automate actions like a human user.
Skyvern	2024	Early-stage agent infrastructure company building autonomous AI agents that can perform web-based tasks via a custom browser and decision engine.
Orby AI	2022	Enterprise automation platform using a proprietary Large Action Model (LAM) to generate no-code workflows by observing user behavior across applications.

Sandbox, Evals, and testing

Company	Founded	Description
Halluminate	2023	AgentOps infrastructure company providing model-driven evaluation tools and testing sandboxes to ensure reliability of AI agents.
Veris AI	2024	Platform for training and validating enterprise AI agents through high-fidelity simulation environments and real-world scenario modeling.
Coval	2023	AI orchestration platform enabling natural language

		interfaces to enterprise tools, allowing users to automate workflows without writing code.
Vocera (aka cekura)	2024	Security-focused AI agent platform designed to autonomously detect, triage, and respond to digital threats in real time across enterprise environments.

Special-purpose (or single use) AI Agents - can do one or two things very well

Company	Founded	Description
Fireflies	2016	AI meeting assistant that records, transcribes, and summarizes conversations to automate note-taking and action item tracking.
Circleback	2024	AI agent that autonomously follows up on meetings and tasks by reviewing transcripts, updating tools, and drafting follow-up messages.
Outter	2024	Workflow automation agent that turns conversations and notes into structured actions across productivity tools like Notion and Asana.
Jasper	2021	AI content generation platform for marketing teams, enabling fast creation of blogs, emails, and campaigns using LLMs.
Writer	2020	Enterprise-grade generative AI platform for content creation, brand voice consistency, and knowledge-driven writing across teams.
Apriora	2024	AI agent platform that automates recruiting workflows—from sourcing and screening to scheduling—through LLM-powered interactions.
Rogo	2023	AI analyst assistant that helps teams ask questions of their data in natural language and generates accurate insights instantly.
Gama	2022	AI-powered presentation and document creation tool that transforms text prompts into polished slide decks and visual content.

Incumbents That Matter (strengths and weaknesses)

As AI agents transition from novelty to enterprise staple, several incumbent companies (large tech firms and established software players) will heavily influence the trajectory. Understanding their strengths and weaknesses helps us see where startups can fit in or compete.

OpenAI and Anthropic	They historically focus on the model itself, not the end-to-end agent solution. However, they are starting to build agent tooling (OpenAI's function calls, Anthropic's constitutional AI for safer agents), but **they won't capture every layer**. It's possible that they will leave room for others.
Microsoft, Salesforce, and Google	They are already deeply integrated in the enterprise and will use their existing products (e.g. office, Google Suite, etc.) as a way to introduce more AI features into workflows. They already have the trust and integration, but they are likely not to take risks while AI hallucinates.
UiPath, Automation Anywhere, and RPA	They also have robust management suites for bots (orchestration, error logging). UiPath, for example, has an extensive library of connectors and a studio for building automation scripts, assets that could be extended with AI. Weaknesses *is that* traditional RPA bots are largely deterministic and brittle when facing variability.
Zapier, IFTTT, and Integration Platforms	These are incumbents in no-code workflow automation for business users. Zapier has 5,000+ app integrations and is used widely by non-engineers to connect apps (e.g., "when a form is submitted on Typeform, create a lead in HubSpot"). *Weaknesses:* These tools historically handle straightforward rule-based flows; complex multi-step reasoning or iterative planning is beyond their scope

What Would Make Us Invest

With so much activity in AI agents, we need high conviction to pick winners. Here are the key factors and signals we look for that would make us eager to invest in a company in this space:

- **Solving the "hard problems"** (reliability, safety, compliance) – The strongest value will accrue to those who figure out the toughest challenges holding back agent adoption. If a startup demonstrates a solution that materially reduces **hallucinations** or **failures** (e.g. an eval method catching 90% of errors before deployment), that's interesting. If there is

baked-in compliance (ensuring agents follow regulations or company policies) is important for the enterprise. For example, if **Halluminate's** model-driven testing can guarantee an agent meets certain accuracy thresholds, or if an **AgentOps (portfolio company)** platform can *prove* it prevents costly mistakes, those are game-changers.

- **"Whole product"** approach and platform stickiness: We prefer companies that can evolve into platforms rather than one-off tools. An ideal investment is a startup that starts with a beachhead (say monitoring) but can expand across the AgentOps stack or build a marketplace. We look for signs of ecosystem creation. For example, a company has third-party developers adding integrations or a community contributing plugins. This often shows up as strong open-source adoption or API usage growth. If developers are building on a platform (LangChain is an example with its integrations), it increases stickiness and network effects. For no-code platforms, we'd want to see **templates**, an "app store" of agents, or other indications that users are sharing and standardizing on it.

- **Bottom-up traction with top-down buy-in:** Given our thesis on adoption, the best companies will capture both bottoms-up usage and enterprise deals. We'd be very excited about metrics like: thousands of weekly active users (developers or ops analysts) using a product organically (bottom-up validation), *and* a growing pipeline of enterprise contracts or pilots (top-down validation). For example, if an agent builder tool starts as a free community edition that individual employees use, and we see it spread to dozens of teams in a Fortune 500, that's a signal. Then, if the CIO of that company approaches to get an enterprise license (because they noticed many employees using it), it shows bottom-up converted to top-down. We'd invest in a startup that has this usage pattern and can articulate a clear path to monetize that interest via enterprise SaaS agreements. Early signs might be open-source GitHub stars and community growth coupled with inbound requests from large logos.

- **Team expertise and speed:** We place weight on teams who deeply understand both AI tech and enterprise needs. Given how fast this field moves, a team that can *rapidly iterate* in response to new model capabilities or customer feedback is crucial. We'd look at the founders' backgrounds; for example, did they build AI systems at scale before (ex-Google, OpenAI, etc.)? Have they surrounded themselves with advisors or early employees from relevant areas (AI research, enterprise sales, etc.)? An ideal founding team might be a combo of a researcher/engineer who built agent systems plus a product lead who scaled a B2B SaaS.

- **Evidence of moats forming:** In AI agents, potential moats include: data network effects (e.g., the more agents run on the platform, the smarter it gets at evaluation or suggesting improvements), community adoption (open-source dominance, developer mindshare), integration depth (becoming so embedded in customer workflows that switching costs are high), and possibly proprietary models or IP (though models tend to commoditize).

- **Economics and scalable model:** viable business model for scaling. A clear idea or direction for pricing (SaaS subscription per agent or per user, usage-based, etc.) and whether early customers are willing to pay. We'd be cautious if a solution requires too much custom work for each customer. Ideally, the startup can show a growing ARR, low churn even at an early stage, indicating that once a company adopts the tool for their agents, they rely on it (stickiness).