Stubber

AI Agents: Adding a Smarter Layer to Work

AI agents are reshaping how work is organized and should be welcomed as an additional layer that streamlines and takes ownership of complete workflows

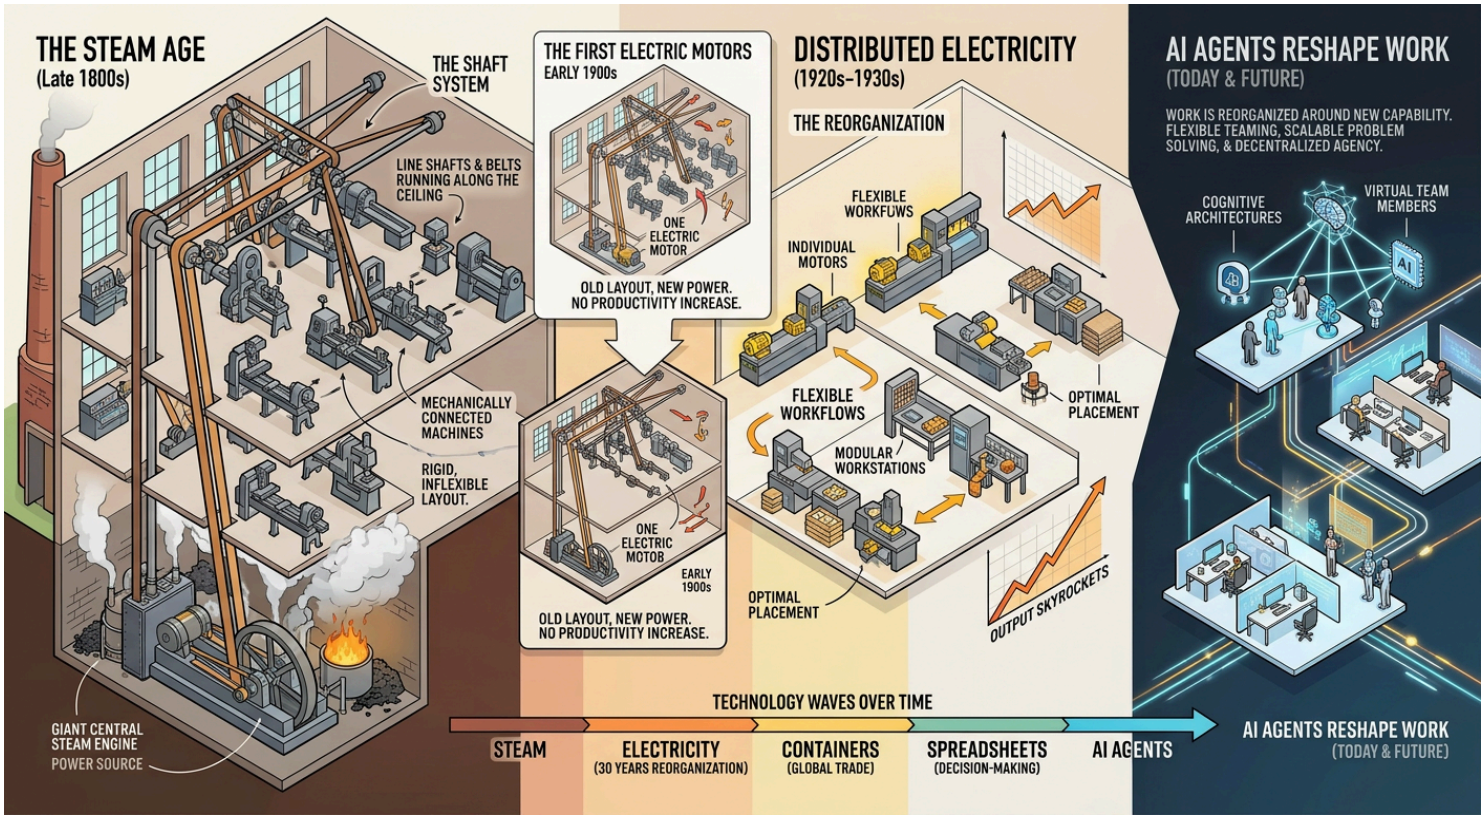

In the late 1800s factories ran on steam engines. A single giant engine in the basement powered the entire building through line shafts and belts running across the ceiling. Every machine connected to this central power source. When electric motors were first introduced, most factories replaced the steam engine with an electric one, but kept the same layout. The result? Almost no productivity improvement.

It took decades before factory owners realized the real advantage of electricity: each machine could have its own motor.

Once factories reorganized around distributed power, everything changed. Production lines became flexible. Machines could be placed where they were most efficient. Workflows became modular and output skyrocketed, although economic historians estimate that it took 30 years before full adoption and the productivity benefits were realized.

Every major technological wave follows the same pattern. At first, companies bolt the new technology onto old systems. However, the real breakthroughs happen when organizations redesign themselves around the new capability.

Electricity reshaped factories. Containers reshaped global trade. Spreadsheets reshaped decision-making. And now, AI agents are beginning to reshape how work itself is organized.

A Framework for Executives

With all the noise around AI – and the growing sense of “AI fatigue” – many executives are asking the same question: Where do we actually start?

A helpful way to think about this technological wave is to break it into a few clear shifts.

1. Legacy AI versus large language models (LLMs)

For many years, AI meant highly specialized models built for very specific tasks: one model for optical character recognition, another for sentiment analysis, another for fraud detection.

Building these systems required large volumes of structured data, specialized expertise and significant computational resources.

Legacy AI was primarily about understanding data: identifying patterns, classifying information and generating predictions.

This form of AI remains extremely valuable today and often plays an important role alongside modern AI systems, providing structured insights and analytics that can feed into more advanced automation.

2. LLM

LLMs represent a fundamentally different approach. Instead of building a different model for every task, LLMs provide a form of general-purpose intelligence that can perform a wide variety of cognitive work.

Trained on vast amounts of language and knowledge, these models reason about problems far more flexibly. This creates a powerful new interface for business systems: language.

Instead of writing complex software or training specialized models, businesses can increasingly describe the task in natural language and allow the system to determine how to complete it.

LLMs also tolerate messy environments remarkably well. They do not require perfectly structured datasets or rigid schemas. Much like a human employee, they can interpret incomplete information, infer missing context and adapt to the realities of how work actually happens inside organizations.

The barrier to entry has dropped dramatically.

3. Human-augmenting AI versus independent AI agents

Most early applications of LLMs have focused on assisting humans. Tools like ChatGPT, Claude and Copilot help employees write, analyze, research and summarize information more quickly. These systems augment human capability and can deliver meaningful productivity gains.

However, this is only the first stage of the transformation.

Using AI purely as an assistant is somewhat like replacing a steam engine in a factory with a single electric motor – the underlying organizational structure remains largely unchanged.

The deeper shift comes with independent AI agents. AI agents do not simply assist humans; they carry out work autonomously.

They can interact directly with customers, staff and internal systems to complete tasks from start to finish.

Examples include agents that operate across revenue, service and operational workflows:

- AI sales executive: conducts needs analysis, generates quotations, verifies information, completes onboarding and closes the sale.
- AI support agent: investigates a customer issue across internal systems, diagnoses the root cause and resolves the problem.
- AI claims agent: captures a claim and coordinates assessors, repair contractors, and internal teams to resolve the case end-to-end.

This distinction is important.

Answering frequently asked questions is not an AI agent. A true agent resolves the problem end-to-end without human intervention.

The strategic opportunity for organizations lies in identifying where these autonomous agents can take ownership of complete workflows rather than simply assisting with individual tasks.

The Three Stages of AI Adoption

1. Insight: legacy AI that analyzes data.
2. Assistance: LLM tools that help employees work faster.
3. Autonomy: AI agents that execute work independently.

The First Steps

For executives, the key is to approach this wave of technology in a structured but pragmatic way.

1. Go deep, not wide

One of the most common mistakes organizations make with new technologies is spreading efforts too thin.

Instead of launching dozens of small AI initiatives across the business, it is far more effective to choose one or two meaningful workflows and go deep. Focus on areas where autonomous agents can own a complete process, typically customer-facing functions such as sales or support.

When an AI agent can handle an entire workflow end-to-end, the efficiency gains become clear very quickly.

2. Allocate a budget for experimentation

In his book, *The Innovator's Dilemma*, Harvard professor Clayton Christensen explains why large organizations often miss disruptive technologies.

Established companies rely heavily on forecasting, return on investment (ROI) models and detailed analytics before committing resources. However, with emerging technologies, those numbers simply do not exist yet.

When a new wave arrives, there is no reliable data, which means traditional decision frameworks fail.

The way successful organizations navigate this is simple: allocate a defined budget for experimentation. Instead of demanding perfect forecasts, leadership sets aside a small pool of capital to explore new capabilities. The goal is not immediate ROI, but learning.

Once successful use cases emerge, investment can scale quickly.

In other words: experiment first, optimize later.

3. Start with a focused AI agent deployment

The most practical way to begin is with a targeted implementation where an AI agent can take ownership of a clearly defined workflow.

Typical starting points include:

- Sales qualification and quoting.
- Customer onboarding.
- Technical support resolution.
- Service scheduling and follow-ups.

These processes are often repetitive, structured and high-volume, making them ideal for autonomous agents.

In early deployments, AI agents are already handling complete support interactions: diagnosing network issues, scheduling technicians and resolving customer problems without human intervention.

Embracing Real Transformation

Every major technological wave follows a familiar pattern. At first, organizations experiment cautiously. The technology is used in small ways – often bolted onto existing processes. The impact is incremental.

However, over time, some companies begin to reorganize themselves around the new capability. They rethink workflows. They redesign roles. They allow the technology to take ownership of entire pieces of work. That's when the real transformation happens.

The question for executives is not whether AI will affect their industry. It is whether their organization will treat AI as a tool that assists employees or as a new operational layer that can execute work independently.

Those that experiment early and redesign their workflows around autonomous agents will unlock extraordinary efficiency gains. Those that wait for perfect certainty may discover that the wave has already passed them by.

How Stubber Can Help Organizations

Platforms are emerging to help organizations deploy these capabilities quickly. One such platform is Stubber.

AI Agent Starter Packages

To make experimentation easy, Stubber offers starter packages designed specifically for organizations exploring agentic AI. These packages include:

- Prepaid AI credits that cover usage across the full stack – large language models, databases, vector stores and communication systems.
- Implementation hours with Stubber's solution engineering team to design and deploy real use cases.
- No long-term commitment, allowing organizations to explore opportunities with minimal risk.

The platform integrates the full AI agent stack – from models to messaging channels – meaning businesses can begin deploying agents without needing to onboard multiple vendors or assemble complex infrastructure. This allows teams to focus on what matters most: identifying the workflows where autonomous agents can create the greatest impact.

Stubber Inc.

Read more on the Stubber blog
AI Agents: The Smart Layer