

# The Path to Autonomous Operations: How Startup CEOs Can Use GenAI for Scale and Competitive Advantage

Published 31 October 2023 - ID G00800292 - 15 min read

By Analyst(s): Mark Petty

Initiatives: [Scalable Business Operations for Tech CEOs](#); [Investing in Readiness to Scale](#); [Generative AI Resource Center](#)

The next billion-dollar SaaS startups will have generative AI embedded into their DNA. Tech CEOs will find that these high-growth startups will achieve scalability, capital efficiency and competitive advantage quicker with GenAI, by transitioning from ad hoc use to fully autonomous operations.

## Overview

### Key Findings

- Future market leaders will quickly gain first-mover advantage in their industries by aggressively adopting generative AI (GenAI) before their competitors. GenAI offers a 12- to 24-month window for tech providers to gain a competitive advantage in their operating models.
- Autonomous operations will require a new multiagent operating model, where each agent will have a defined role in a virtual organization chart that mimics the hierarchical structure of enterprises, mirroring real-world team dynamics.

### Recommendations

To scale business operations with GenAI, tech CEOs must:

- Keep ahead of competition by building a roadmap to autonomous operations, progressing through five maturity levels. Plan to transition from initial ad hoc use of GenAI to augmented departments within 24 months. Envision that by 36 months, GenAI will be integrated companywide.

- Take care to adopt GenAI agents that operate within a functional multiagent operating model, which will be governed by roles, principles, procedures, decision-making rights and process responsibilities.

## Strategic Planning Assumptions

- By 2025, GenAI will enable startups to improve revenue per employee by 25%.
- By 2028, GenAI native SaaS startups will need one-tenth of current funding to generate the same revenue as today.
- By 2029, 50% of high-growth startups will have adopted a multiagent operating model, up from less than 1% today.

## Analysis

GenAI-native businesses will outcompete legacy organizations within five years, reimaging every operational process that can be augmented or automated by GenAI. This presents an opportunity for tech CEOs to get to scale with fewer resources.

### Productivity Gains

- Tech CEOs report they are already exploring how they will use GenAI within their operations to boost productivity. <sup>1</sup> Many believe it is startups that will benefit the most from the early wave of GenAI adoption, that they can accomplish more with fewer employees. Investors are also beginning to focus their capital on companies that are driving exponential value with GenAI in their operating models.
- GenAI offers startups capabilities previously only accessible to tech giants – advanced analytics and automation to name a few. This emerging set of capabilities is being used by startups to transform their operations and will bring about new modes of operating at scale.

### Competitive Advantage

- GenAI offers a short-term advantage comparable to the cloud's early days. Tech CEOs that utilize GenAI to augment operations have the opportunity to not only improve incrementally, but also develop a sustainable lead over competitors before using generative AI in this way becomes table-stakes.

- GenAI offers an immediate 12- to 24-month window for startups to gain a competitive advantage through productivity gains. In the next three to five years, the productivity gains will extend into areas of autonomous agents, which will enable systems that can build processes, workflows and trigger actions based on instructions given. Employees will become designers and orchestrators of systems and outputs.
- To maintain a competitive advantage into 2028 and beyond, startups must have planned in advance and adopted a multiagent operating model.

Here is a glimpse of this future:

## **It's 2028. Today's GenAI-native tech startups are fueled by hybrid human-AI teams.**

Incumbents lagging in AI adoption have been disrupted by nimble, rapidly growing, AI-native startups — and are scrambling to catch up.

Autonomous generative AI agents handle thousands of complex business processes in parallel, improving themselves continuously through reinforcement learning and human feedback.

Specialized agents deconstruct complex problems into small chunks. Each agent has a defined role in a virtual organization chart. Jobs are broken down into modular steps handled by specialized roles, mirroring real-world team dynamics to prevent unproductive loops and cascading failures.

“Product managers” analyze markets. “Engineers” build systems. “Marketers” drive conversions. “Customer support” automate responses, boosting satisfaction. “Financial planners” project future investment ROI. “HR” source top-tier candidates and onboard talent to culture. “Sellers” research prospects and deliver personalized outreach. Each agent has a defined role aligned to human domain expertise.

Humans provide oversight to augment capabilities and guide autonomous agents.

Standard operating procedures (SOPs) encode best practices into the agents, guiding their structured collaboration.

“Project managers” coordinate complex workflows between interdependent agents facilitating knowledge handoffs. “Executives” analyze agent output and course-correct actions. Shared data assets and API integrations fuel insights, giving agents domain knowledge and context to make decisions.

A fractal ecosystem of third-party multiagent operating models has emerged and marketplaces offer on-demand access to agents that are fine-tuned to deliver specialist domain expertise, algorithms and proprietary data.

Products are autocomposed and adapted to the needs of individual customers. Infrastructure resources are autoscaled.

Productivity is accelerated tenfold over legacy businesses.

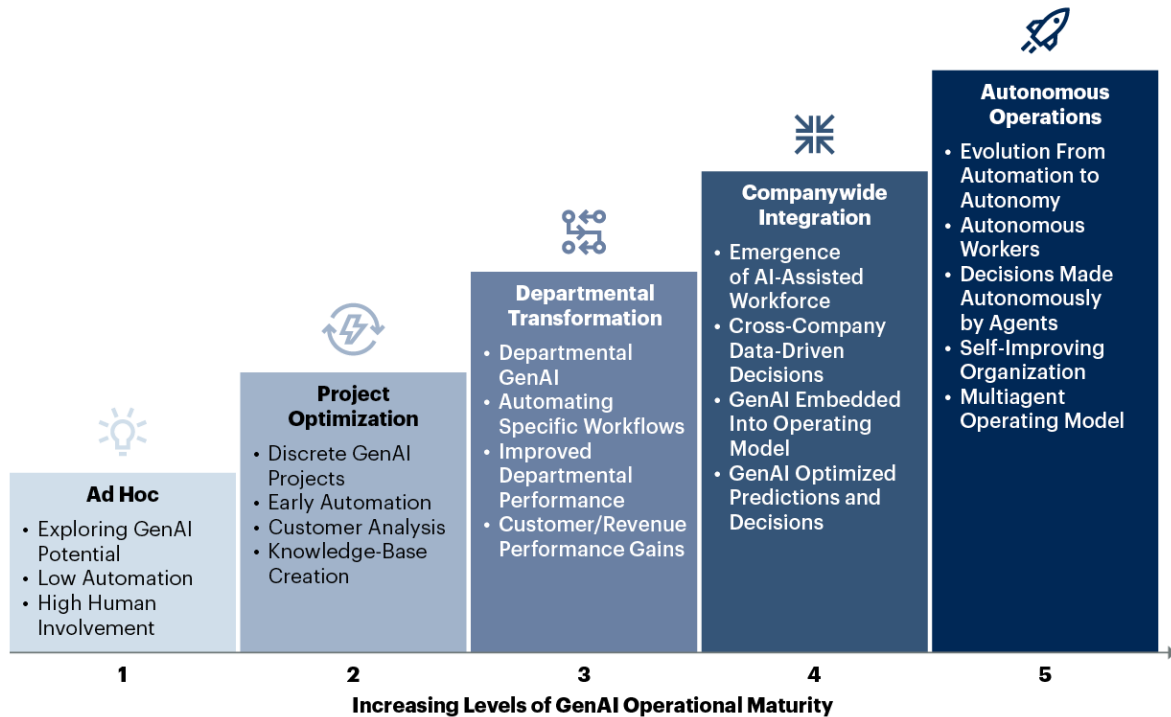
To realize this vision and emerge as a frontrunner by 2028, tech CEOs can take a stepped approach to absorbing GenAI across five levels of maturity – from foundational use up to autonomous operations. From documenting processes and procedures, gaining predictive insights from data, automating processes, optimizing decisions and simulating scenarios, to building full GenAI-powered autonomous operations.

## The GenAI Operations Scalability Matrix

As startups adopt GenAI to enhance operational efficiency, they will advance through five stages of maturity (see Figure 1). Gartner refers to this as the “GenAI Operations Scalability Matrix.” Tech CEOs can pursue exponential performance gains as the technology improves over time and will be able to accomplish more with less at each maturity level.

Figure 1. GenAI Operations Scalability Matrix

**GenAI Operations Scalability Matrix**



Source: Gartner

AI = artificial intelligence; GenAI = generative artificial intelligence

800292\_C

The following sections describe each of these phases.

**Level 1 – Ad Hoc Use**

- Use GenAI to document processes and workflows.
- Have GenAI applications like OpenAI’s ChatGPT, Anthropic’s Claude and Google Bard turn processes into standard operating procedures (SOPs) and guidelines.
- Leverage GenAI applications in an ad hoc manner across daily tasks and workflows.
- Concentrate on leveraging GenAI to save time through semiautomated creation of templates, checklists and other internal operational documents.

**Level 2 – Project Optimization**

- Leverage GenAI tools to analyze data for discrete AI projects (such as customer sentiment analysis) to help inform product roadmaps.

- Adopt consistent top practice approaches for prompt engineering. A well-constructed prompt takes advantage of GenAI reasoning abilities while shaping them to meet specific enterprise requirements.
- Augment specific enterprise data, either by passing data or integrating search engines to locate and retrieve any information necessary to construct a prompt, including proprietary information.
- Leverage “in-context learning” to teach large language models (LLMs) how to complete a task in a particular way.
- Pick a self-contained process or subsystem to automate first, such as prioritizing customer support. Use rapid iterations and user feedback to improve the automation. GenAI can accelerate building minimum viable product (MVP) bots, for example, handling common customer inquiries.
- Document repetitive tasks that could be candidates for automation.

### Level 3 – Departmental Transformation

- Microsoft and Google have both announced the embedding of GenAI in their core offerings. As have other department-specific tools, such as CRMs like Salesforce and HubSpot. Leveraging these offerings (when available) will deliver specific departmental gains. However, for GenAI to truly be transformative to the enterprise, GenAI must have access to current, proprietary and contextual information.
- Enable departments like sales, marketing and support with AI content creation tools to accelerate productivity and use “chain of thought” prompting. This is in order to enable a LLM to break down complex problems into multiple steps, each providing an intermediary solution that feeds into the next step. Open-source projects are emerging (such as LangChain to manage these chains), along with integrations into other tools, databases and platforms.
- Automate workflows within departments aimed at discrete generating efficiency gains.
- Task-specific enterprise assistants for business functions like CRM, HR, accounts receivable and others. Plug-ins such as Zapier for ChatGPT can assist with integrations into other systems.
- Leverage GenAI for predictive analytics at the department level, like forecasting KPIs, customer lifetime value and churn risk.

- Build high-fidelity customer service bots that answer contextual questions. Take documents and chunk them down into small elements in a vector database.
- Things at this level are tactical arrangements of prompt engineering.

## Level 4 – Companywide Integration

- Automate interactions and handoffs between departments, break down data and process silos through AI coordination to facilitate cross-company, data-driven decisions.
- Build agents that can iteratively and autonomously call LLMs, plus other models and tools to accomplish a multistep task. Autonomous agents are combined systems that achieve defined goals without human intervention. Tasks that do not follow a fixed workflow will be enabled at this level through a flexible approach to calling LLMs and different tools that deepens context. Agents that can recursively access LLMs, search enterprise data and other tools, while persisting context and previous interactions will accomplish this. Tools that can spin up customized generative AI agents are being developed and are emerging fast. AutoGPT and Council are examples of GenAI agent platforms.
- Develop companywide data taxonomy and knowledge graphs. The creation of knowledge bases of curated enterprise content will be essential. Create a common repository as a staging area for initial collection and processing of information from enterprise systems.
- Use foundation model fine-tuning to adapt a LLM foundation model with the enterprises's own knowledge in order to perform new tasks or respond to new patterns. Fine-tuning is not currently possible for some of the most advanced proprietary LLM models, but it is possible for open-source models.
- Provide teams with an AI assistant advisor for real-time data insights and continuously assess decisions against goals and suggest improvements. GenAI-optimized predictions and decisions will be made possible with tools like ChatGPT code interpreter.

## Level 5 – Autonomous Operations

- Hybrid human-AI teams become the norm, with AI “teammates.”

- Generative interfaces give access to insights in real time — that previously would have required analysts. With predictive models and simulations, tech CEOs can gain a competitive edge that is usually only available to larger companies.
- Autonomous operations can be realized, first in part, then with full workflows automated and optimized by generative AI agents, which will progressively shrink the overhead needed to support increased levels of revenue.
- Multiagent generative system designs will enable multiple agents (discriminative and generative models) to collectively cooperate to deconstruct and solve complex problems.
- Each agent has a defined role in a virtual organization chart that mimics the hierarchical structure of enterprises. Jobs are broken down into modular steps handled by specialized roles, mirroring real-world team dynamics to prevent unproductive loops, and minimizing cascading failures from hallucination.
- Systems can scale intelligently by adding more agents for increased workloads.
- Building a GenAI-first operating model could enable a scalable “self-tuning” organization, agents can handle hundreds of complex business processes in parallel, improving themselves continuously through reinforcement learning to improve processes automatically. Adding self-tuning and analytics capabilities creates a startup where both systems and people are continuously improving.

## Plan for an Autonomous Operations Future

GenAI is moving quickly through the Hype Cycle (see [Video: Generative AI for Real-World Business Productivity](#)). The rate of adoption, innovation and progress has rarely been seen before. Tech CEOs must make plans to stay ahead of the technologies as they evolve.

### Get Ahead of the Hype

- Early in the life cycle of an emerging technology, most of the market isn't aware of the technology and doesn't understand the impact. It's largely unknown, and as a result, first movers have the opportunity to gain an advantage — both when applying the technology and as providers.

- Full autonomous operations will be enabled by the maturing of several distinct technologies in the domain of GenAI including:
  - Autonomous agents.
  - Reinforcement learning from human feedback.
  - Domain-specific GenAI models.
  - Language operations (LangOps) and GenAI application orchestration frameworks.

(See [Hype Cycle for Generative AI, 2023](#).)

## Plan for Competitive Advantage

- There are clear parallels between GenAI today and the rise of cloud computing over the past decade – in terms of providing tech CEOs with a window of competitive advantage. Those who aggressively adopt GenAI will more likely become future market leaders, but these capabilities will eventually become widespread and standard across industries.
- Plan for a 12- to 24-month window to leverage available GenAI capabilities for efficiency and growth before it becomes commonplace. Expect further productivity gains as GenAI extends into areas of autonomous agents, which will afford a further two to three years of competitive advantage. Tech CEOs should prioritize building their internal AI skills, infrastructure and workflows in order to maximize this advantage.

## Prepare for a Multiagent Operating Model

- To build fully functioning autonomous operations, care must be taken for designing agents to function within the new operating model, where each agent will have a defined role in a virtual organization chart that mimics the hierarchical structure of enterprises. These will mirror real-world team dynamics to minimize cascading failures from hallucination.
- Appropriate governance through SOPs and decision rights will be necessary. Research is showing that a role-based approach generates more complete solutions versus chat-based agents.<sup>2</sup> Tech CEOs can now start to identify capability gaps and governance weaknesses as they scale.

## Document and Define Processes

- Just as time is taken to justify, hire, onboard, train and assimilate an employee into a company, so too will this same level of care be required for pools of agents that must work together under constraints to work productively for the company's operations.
- For a scaling startup, care must be given to document processes, procedures and guidelines for employees – still, mistakes happen and course correction, guidance and training is necessary. These are the expectations placed on every member of the team, including the CEO. Expect and prepare for similar needs and methodologies to be applied to AI workers.
- Governance over autonomous agents' decision making will require behavioral guardrails including roles, principles, SOPs, decision-making rights and process responsibilities.
- Tech CEOs can establish a head start now by building out policies, principles, procedures and practices (see [Startup Lift: Establish the Practices and Procedures Needed to Get to Product-Market Fit](#)).

## Prepare the Company by Becoming a Data-Driven Leader

- In an autonomous operations future, leadership will transition from one of people management to orchestration through data and feedback. Some of the fastest-growing tech companies have grown through high levels of autonomy and delegated responsibility through the organization. It is this exact posture that will be required for autonomous operations adoption.
- Companies with data-driven leaders have cultures that base decisions on facts and data instead of "gut feelings." This results in better delegation from leadership and autonomy from employees. To capitalize on the future, companies must adopt this posture (see [Tech CEO Insight: Develop These 4 Attributes to Become a Data-Driven Leader](#)).

## Prepare for an Autonomous Business Future

- When combined with other emerging tech-enabled trends – including smart products, machine customers and programmable economy – autonomous operations will enable the transition to the next distinct tech-enabled, macro wave era of business, "autonomous business."

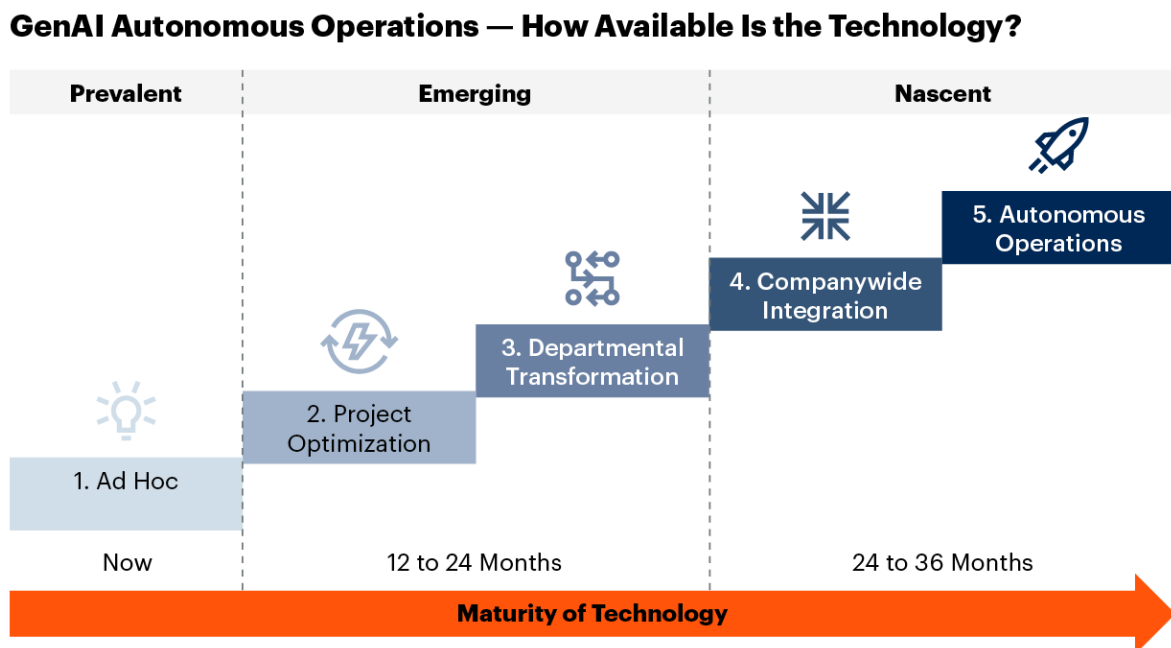
- “Autonomous business” is Gartner’s holding term for this shift and is defined as “A style of business partly governed and majority operated by self-learning software agents that provides smart products and services to machine-customer-prevalent markets, operating in a programmable economy.” See [Autonomous Business Is the Next Tech-Enabled Strategic Growth Curve for Pioneer Enterprises](#).

## Background and Context

### Technical Challenges

Tech CEOs can pursue exponential performance gains as technology improves over time and will be able to accomplish more with less at each maturity level. Level 1 establishes basic foundations, while Level 5 enables exponential gains. Figure 2 shows the availability and maturity of the technology to support each level.

**Figure 2. GenAI Autonomous Operations — How Available Is the Technology?**



Source: Gartner  
 GenAI = generative artificial intelligence  
 800292\_C

Existing LLM-based agents primarily focus on solving simple dialogue tasks, complex tasks are hard to scale — mainly due to the LLM hallucination problem. Hallucinations become cascading when chaining multiple intelligent generative autonomous agents, resulting in a failure to effectively address complex real-world problems at scale.

Technical challenges that must be overcome for autonomous operations to be realized include:

1. **Brittle workflows.** Standardized workflows are difficult as exceptions and edge cases are hard to adapt to. Coordinating workflows remains an open research problem.
2. **Coordination.** Effective mechanisms are needed for agents to coordinate and share knowledge.
3. **Lack of feedback loops.** No systematic mechanisms exist for agents to correct themselves based on results.
4. **Reasoning limitations.** While LLMs display an incredible aptitude for reasoning, LLMs today have limited capabilities for complex logical reasoning, causal inferences and anticipating complex chains of events. This constrains their ability to make rational decisions.
5. **Scaling costs.** Scaling the compute needed for running thousands of AI agents simultaneously for an enterprise will have huge demands on computational resources which will prove costly.
6. **Incomplete domain knowledge.** While LLMs have made great progress, they still lack the depth of specialized domain knowledge that comes from human experience and education. Capturing this remains challenging.
7. **Lack of common sense.** Autonomous generative agents today act in literal ways and without broader human context and common sense, this can result in suboptimal (even absurd) behaviors.
8. **Hallucinations.** Existing LLM-based autonomous generative agents work well for simple tasks but struggle with more complex ones, primarily due to cascading errors from hallucinations.

Technologies to support fully autonomous operations are nascent, but open-source projects such as MetaGPT are emerging to orchestrate LLM-based agents. MetaGPT defines specialized agent roles and actions specific to the software development life cycle to generate code. Research shows that this role-based approach generates more complete solutions versus chat-based agents. <sup>2</sup>

## Acronym Key and Glossary Terms

AI	artificial intelligence
API	application programming interface
CEO	chief executive officer
ChatGPT	Chat Generative Pretrained Transformer
CRM	customer relationship management
GenAI	generative artificial intelligence
HR	human resources
KPI	key performance indicator
LangOps	language operations
LLM	large language model
MVP	minimum viable product
ROI	return on investment
SaaS	software as a service
SOP	standard operating procedure

## Evidence

<sup>1</sup> A social media analysis of 9,788 conversations showed that startups are using GenAI as a tool to enhance productivity. Gartner conducts social listening analysis leveraging third-party data tools. Due to its qualitative and organic nature, the results should not be used separately from the rest of this research. No conclusions should be drawn from this data alone. Social media data in reference is from 1 July 2022 to 31 August 2023 in all geographies (except China) and recognized languages.

<sup>2</sup> Sirui Hong, Xiawu Zheng, Jonathan Chen, Yuheng Cheng, Jinlin Wang, Ceyao Zhang, Zili Wang, Steven Ka Shing Yau, Zijuan Lin, Liyang Zhou, Chenyu Ran, Lingfeng Xiao, Chenglin Wu, [“Meta Programming for Multi-Agent Collaborative Framework,”](#) arXiv, 17 August 2023.

## Recommended by the Author

Some documents may not be available as part of your current Gartner subscription.

[Top GenAI Operational Use Cases and Risks on the Minds of Startup Tech CEOs](#)

[GenAI Funding Opportunities for Tech CEOs: Fast and Fleeting](#)

[Video: Generative AI for Real-World Business Productivity](#)

---

© 2023 Gartner, Inc. and/or its affiliates. All rights reserved. Gartner is a registered trademark of Gartner, Inc. and its affiliates. This publication may not be reproduced or distributed in any form without Gartner's prior written permission. It consists of the opinions of Gartner's research organization, which should not be construed as statements of fact. While the information contained in this publication has been obtained from sources believed to be reliable, Gartner disclaims all warranties as to the accuracy, completeness or adequacy of such information. Although Gartner research may address legal and financial issues, Gartner does not provide legal or investment advice and its research should not be construed or used as such. Your access and use of this publication are governed by [Gartner's Usage Policy](#). Gartner prides itself on its reputation for independence and objectivity. Its research is produced independently by its research organization without input or influence from any third party. For further information, see "[Guiding Principles on Independence and Objectivity](#)." Gartner research may not be used as input into or for the training or development of generative artificial intelligence, machine learning, algorithms, software, or related technologies.