Galson.TPI
Technology Prioritization Index

Generative AI

You have already adopted it.
Now make it pay.

TPI Navigator Rank: **Emerging**

TPI Score: **13/21**

Authors:
Christopher Richardson, Principal Analyst
Aliaksei Kharkin, Analyst

01

Everyone is in. Few are paid back.

Between 65 and 72 percent of organizations now run at least one generative AI workload in production. About 60 percent of them report cost overruns the moment they try to scale past a pilot. Adoption was the easy part. The payback is the hard part.

If you run technology at an organization, the question has already moved on. It is no longer whether to adopt generative AI. It is already in your customer service queue, your marketing team, and your developers' editors. The real question is whether you can make it pay, keep it compliant, and trust what it produces.

That is what this report measures. We ran generative AI through Galson's Technology Prioritization Index, the scoring system we apply to every domain we cover, and we tell you what the score should change about your next move.

Up front. Generative AI scores 13 out of 21. We rank it Emerging: further along than most AI categories, genuinely useful in production, and still losing money for most organizations that scale it without discipline.

02

From novelty to infrastructure

Generative AI has crossed from experiment to infrastructure. In early 2026, 72 percent of organizations reported at least one AI workload in production, and 65 percent said they were using generative AI in at least one business function. That is well past pilot territory.

Two forces pushed it there: the tools got easier to reach, and the pressure to turn AI experiments into measurable productivity got heavier. Organizations are folding it into customer service, IT operations, marketing, knowledge work, and software development. How well those deployments actually work still swings hard by industry and use case.

The next round of value will come less from access to powerful models, which everyone now has, and more from disciplined implementation, clean data practices, and connecting AI output to workflows people can rely on.

03

How the TPI works

The TPI is Galson's scoring system for technology domains. We score each domain from 1 (weak) to 3 (strong) on seven factors that decide whether a technology can carry real business weight, not just demo well. The factors add up to a number out of 21 and a Navigator Rank that tells you how to treat it.

Generative AI lands at 13 of 21: Emerging. Here is how it scored, and what each score should change about your plan.

Cases

2/3 Proven Use Cases

Generative AI has real traction in a few places, led by automated content generation and code autocompletion. Newer uses are showing measurable value too: synthetic data creation for ML training, now used by 33 percent of healthcare firms, and customer service chatbots that handle 30 percent of tier-1 inquiries. The ceiling shows up on the hard work. For organization-critical jobs like end-to-end supply chain optimization or financial risk modeling, there are no standard frameworks yet, and only about 10 percent of organizations get repeatable success in complex workflows.

1/3 Cost-Effectiveness & ROI

This is the weakest score, and it is where budgets break. API-based tools are cheap for small tasks. Scale past the pilot and the costs pile up: 60 percent of organizations report cost overruns, driven by infrastructure, scarce specialist talent, and compliance overhead. Even where productivity climbs, the return stays speculative in most pilots, and 54 percent of organizations point to licenses they are not using and tools that duplicate each other.

2/3 Ease of Use & Integration

Getting started has never been faster. Pre-trained foundation models and low-code platforms have cut initial deployment timelines by 40 percent against 2025. The harder deployments still need real expertise: 70 percent of organizations report delays wiring generative AI into legacy data pipelines, and 65 percent cite higher MLOps costs to productionize guardrails like bias mitigation and output validation. Newer capabilities like agentic AI and synthetic data are climbing too, now in use at just over 40 percent of organizations.

1/3 Security & Compliance

The risk is real and the rules are moving. In 2025, 31 percent of organizations reported hallucinations affecting actual decisions, and 52 percent still lack the controls to stop biased or copyrighted data from being ingested. On top of that, 2026 brought a wave of U.S. mandates: California's SB 53, AB 2013, and AB 853, a December 2025 federal Executive Order, and the Colorado AI Act, which together set new requirements for frontier AI safety, training-data transparency, and content disclosure. These are not optional, and most compliance programs are behind.

1/3 Maturity & Ecosystem

The tooling underneath is not stable yet. 65 percent of organizations lean on three or more vendors to cover their needs, and 80 percent report gaps in long-term toolchain compatibility. Open-source models are gaining ground for specific uses but stay niche in production. Standard lifecycle practices like versioning and audit trails are missing from 73 percent of deployments, and 68 percent of IT leaders name thin documentation or community support as a real barrier.

3/3 Market Sentiment

Money and attention are pouring in. The global generative AI market has run past the earlier 41.53 percent CAGR projection to $356 billion. Forecasts for 2026 land between $55 billion and $161 billion, with growth rates of 33 to 47 percent pointing toward $900 billion to $1.5 trillion by 2032 to 2035, pulled by integrations in healthcare, finance, and product development. Read the enthusiasm as fuel, then read the ROI and compliance scores for why it does not turn into value on its own.

3/3 Adoption Growth

Adoption has surged. Between 65 and 72 percent of organizations run at least one generative AI workload in production, and 38 percent of knowledge workers use it daily. Microsoft Copilot reached 41 percent adoption among large organizations by Q1 2026. The catch is depth. Only a fraction have taken it organization-wide, held back by talent shortages, with 45 percent reporting fewer than 10 in-house experts, and by governance gaps that keep accumulating.

04

A tale of two worlds

Generative AI improves almost daily, and people enjoy the new products. The trouble starts when an organization looks for uses beyond the niche ones that justify spending at scale. The category lives in two worlds at once.

High

Proven Use Cases & Maturity

More than 90 percent of large organizations plan to invest in generative AI this year, and the major SaaS and IaaS providers are building it into their products. That sentiment is deceiving. More than 65 percent of poll respondents hold a negative view of generative AI, citing distrust of the output, safety and security worries, the impersonal feel of the results, and fear of job loss. Any organization putting generative AI in front of customers should design for those reactions or risk losing the customer.

Low

Security, cost, and governance

Generative AI's hardest problems sit around data. Public large language models are the most common way in, though not the only one, and choosing between a hosted model and one you run yourself comes down to the price and availability of compute at scale, which can get expensive fast. A Retrieval Augmented Generation setup can cut some of the risk, though building and maintaining one may cost more than it returns. And because there is still little standardized AI regulation, governance falls mostly to individual organizations, while intellectual property disputes over AI-generated work continue to move through the courts.

Nine questions to answer before you scale

The organizations that get a return on generative AI tend to answer these before they expand a pilot, not after

- **Strategy and value.** Does this serve the organization's strategy, and does the use case still justify the full cost of building, running, and governing it?
- **Use cases and risk tiering.** Have you identified the use cases and tiered them by risk, so low-risk tasks move fast while high-risk ones get stricter review?
- **Data and privacy.** Have you defined which data can never go into public AI tools, and are you ready for the privacy, confidentiality, and data-handling obligations that follow?
- **Governance and oversight.** Is there a governance plan with clear human oversight for customer-facing, legal, financial, HR, and operational decisions?
- **Security and model risk.** Are you ready for prompt injection, data leakage, hallucinations, and bias, with logging, validation, and an incident response plan?
- **Regulation and compliance.** Have you mapped your AI use to the privacy, security, copyright, and sector rules that apply, including the new state and industry ones?
- **Vendor and resilience.** Do you know your integrators, consultants, and vendors, and do you have a fallback if one changes pricing, terms, or capability?
- **Change management and readiness.** Is there a real change-management plan, with training, adoption support, and a learning path for employees?
- **Output integrity.** Do you have a process to check AI output for accuracy, bias, and copyright exposure before it drives a business decision?

Answer these before you fund the next proof of concept. Slow or missing answers are the signal to slow down.

Need help turning those answers into a plan? **Galson runs a half-day seminar and strategy program at galson.com/researchlab**

Summary

Bottom line

Generative AI has proven it can shift an industry. The realized value, the kind that shows up on the P&L, has not arrived for most organizations yet. The Emerging score of 13 out of 21 captures that split: real adoption, real capability, and returns that still depend on how well you implement rather than which model you buy.

A lot of investment is landing this year, so expect the tooling, the guardrails, and the proof points to improve. For now, the organizations that win with generative AI are the disciplined ones: tight use cases, clean data, human oversight where it counts, and a clear line from AI output to a workflow people trust. Access to a powerful model is no longer the differentiator. What you do with it is

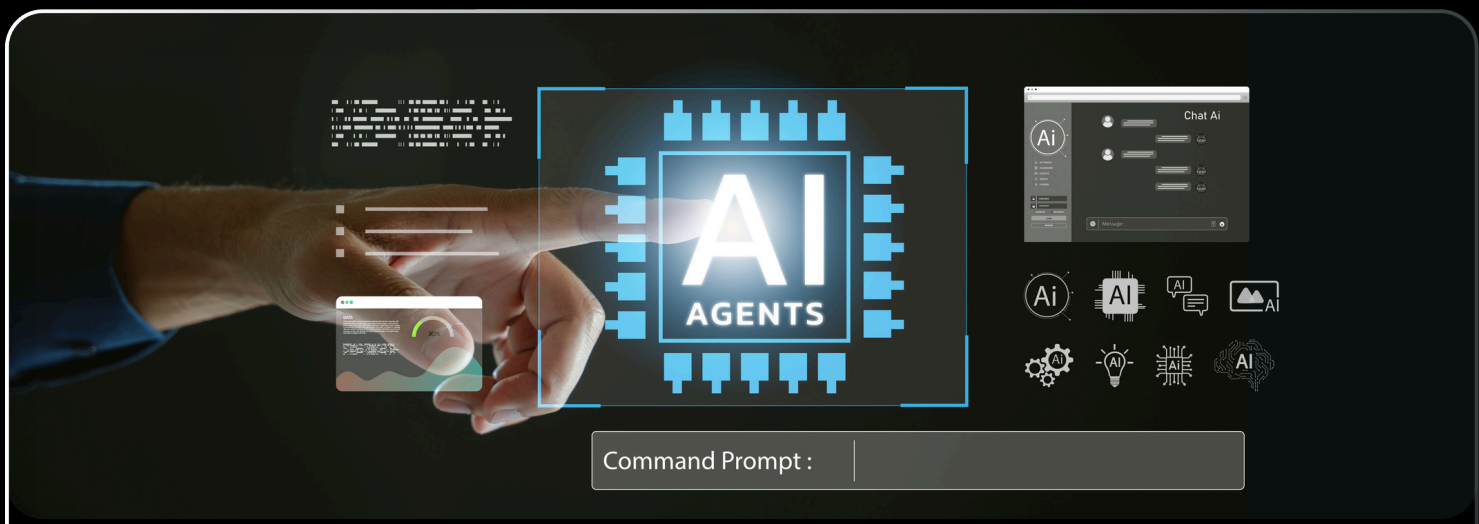

Galson updates this scorecard in real time. The next generative AI update publishes in **Fall 2026.** For a full domain TPI report, or an industry-specific one, visit galson.com or email **hello@galson.com** with the subject TPI Inquiry.