EBOOK

5 best practices that set your software company up for generative AI success

Laying the groundwork to harness the potential of generative AI

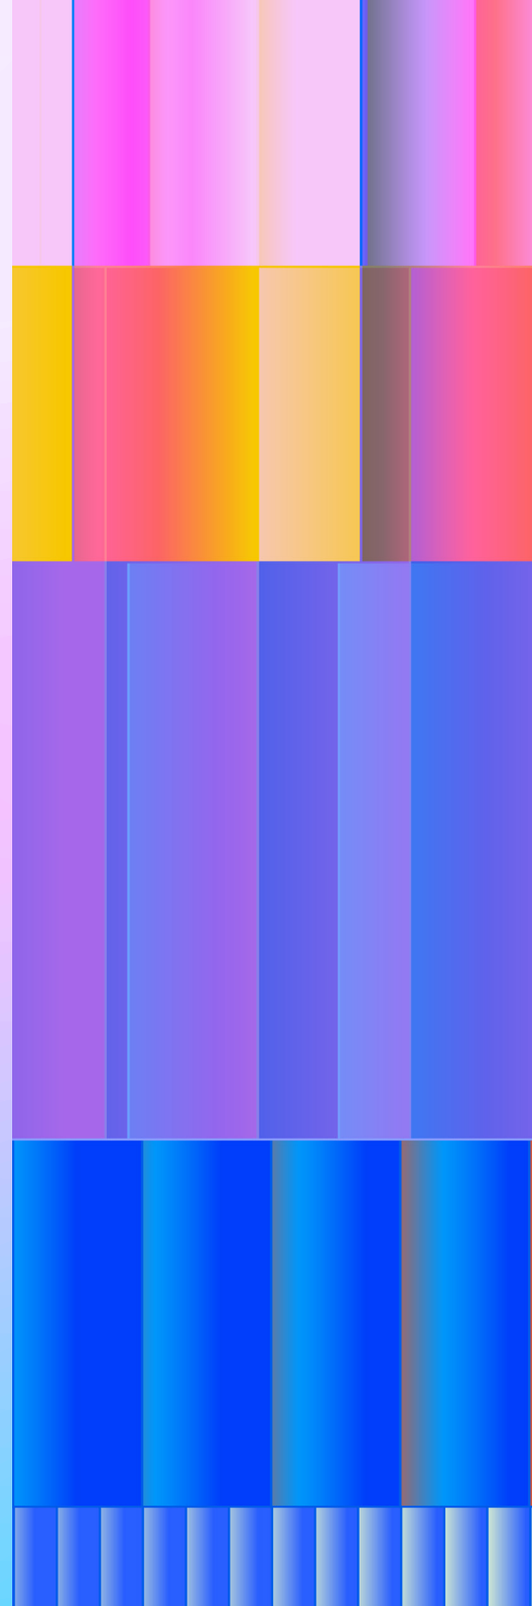

5 best practices that set your software company up for generative AI success

Generative AI is changing everything. Unfortunately, it can be challenging for software companies to make the big changes that are necessary, while keeping their business running. If done right, however, it can pay big dividends – innovation, rapid ideation, developer productivity, speed to market, operational efficiency, and cost savings to name a few. The key is in laying the groundwork to ensure your company can harness generative AI for maximum benefit.

This ebook highlights five best practices identified by software leaders to start effectively using generative AI to drive your business forward:

1. Institute organizational changes

Generative AI has the potential to change everything. To benefit from it, everyone in your organization is going to have to think and operate differently. Starting at the top, executives will need to set a generative AI vision and strategy that redefines what success looks like for the business. This will allow all product and business teams to evaluate and adjust their priorities, resources, and investments to meet these new business goals.

To streamline adoption, employees will need to be educated on why and how their work will be changing, as well as undergo training to gain the knowledge and skills required to effectively harness generative AI. Because generative AI has so many applications that can be infused into almost every aspect of the business, it is not just engineers who will have to change the way they do things – everyone will need to be conversant with generative AI to be more effective.

70%

of the skills used in most jobs will change by 2030[i]

World Economic Forum

80%

of C-suite executives believe AI will kickstart a culture shift where teams are more innovative[i]

World Economic Forum

> "I think generative AI has actually bridged the gap between coders and non-coders, so many teams can actually be involved in generative AI projects."

Naveen Rao,
VP of Generative AI, Databricks[ii]

databricks

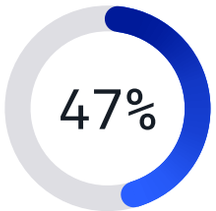

of executives say their people lack the knowledge and skills to effectively implement and scale AI across the enterprise[iii]

IBM Report

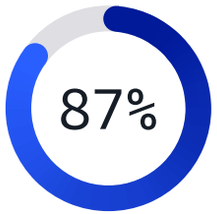

of executives expect jobs to be augmented rather than replaced by generative AI[iii]

IBM Report

Key steps:

1. **Develop a topline generative AI vision and strategy** that helps everyone understand and work towards common objectives.
2. **Create a change management plan** that identifies the projects, processes, and people that will be impacted to enable this new pace of innovation.
3. **Establish the data and technology infrastructure** with the flexibility, scale, and performance required to build, run, and scale generative AI applications in the cloud, on premises, and at the edge. (Learn more: https://aws.amazon.com/ai/generative-ai/)
4. **Revamp workforce planning strategies** to educate and upskill staff, as well as adjust hiring criteria to meet future needs. Users will need to be actively coached to help them understand the potential ramifications of working with these tools, paired with a set of rules or code of conduct that enable users to know what they can and cannot do.

2. Identify ideal use cases

It is important to look at generative AI holistically, not just through one product or one project, to ensure it can have the transformative impact you are looking for on your business. Establishing a top-down generative AI vision and strategy will help all teams, both engineering and business, focus on the best ways to apply generative AI to meet business objectives.

Generative AI is expected to create $900B in productivity uplift with approximately 75% of the total value coming from using generative AI in four functions – customer operations, marketing and sales, software engineering, and R&D.[iv]

Organizations regularly using generative AI, by function, percentage or respondents

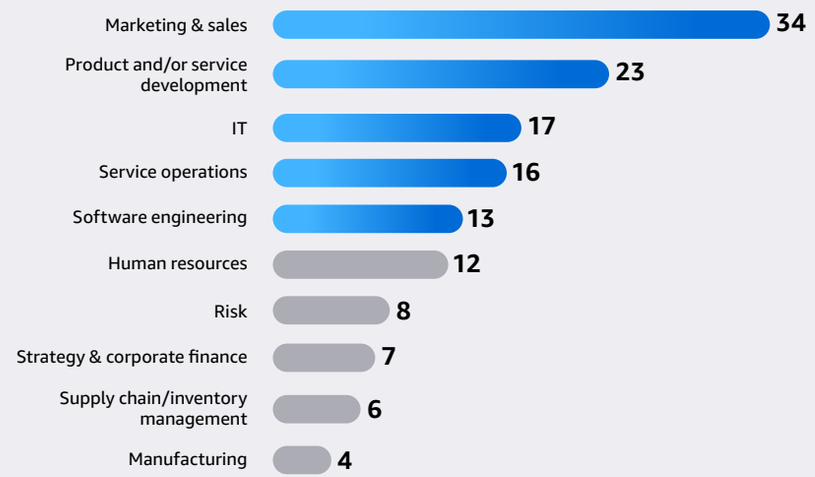

Source: McKinsey Global State of AI Survey, 2024

Each function should look to identify the use cases where generative AI can be applied to derive the biggest benefit. An IDC study in 2024 found increasing operational efficiency, cost savings, profits, and productivity were the outcomes most respondents were looking for from their use of generative AI.

What outcomes are you trying to achieve by leveraging generative AI?

Source: IDC SaaSPath Global Tech Buyer Survey 2024 n=2,825

Once the use cases are identified, whether it's accelerating coding, detecting bugs, creating new services, improving forecasting, reducing risks, etc., they should be clearly defined and scoped. There is a lot of planning that goes into every generative AI implementation. For instance, where will you get the data, what model(s) are you going to use, what are the metrics for success, among other criteria? There's a lot that needs to be figured out and dug into to determine what makes the most sense for your use case and will generate the best results.

Key steps:

1. **Develop a top line generative AI vision and strategy** that helps teams identify the use cases to implement to best address new business objectives. Common generative AI use cases can be found here: https://aws.amazon.com/ai/generative-ai/use-cases/
2. **Scope the use case,** including identifying the data and model(s) that are best for the scenario.
3. **Establish metrics** to measure the progress and levels of success.
4. **Develop a proof of concept (POC)** and good evaluation test case to start to figure out how to get the data, query it, and get the outputs needed for the business.

Rolls out new service that improves customer satisfaction

Salesforce knew revenue leakage could impact up to 10% of a telco's revenue, translating to nearly a hundred billion in annual lost revenue across the industry. To help their telco customers address this issue, Salesforce designed a solution that leverages LLMs to provide business insights from their data, enabling queries, such as "what is the revenue leaks $ impact?", "Which systems drop the most revenue?", and "What are the ways to fix that?" to reduce revenue leakage.

Learn more:
https://www.youtube.com/watch?v=NI1BhsMcG-A

3. Establish good governance

Because new generative AI tools and technology are coming out daily and almost everyone across your organization will be using them in some form or fashion, it is important to ensure that usage is safe and appropriate. Whether teams are scoping a new use case or simply consuming a new service, it is important they understand what is and is not allowable or acceptable.

Establishing guiding AI principles and governance will help your generative AI usage comply with corporate policies and industry regulations and avoid data bias, privacy violations, and other unintended consequences. An AI governance framework allows everyone to work to the same standards with clear guidelines and even legal frameworks that ensure the responsible development and use of generative AI to power your business.

7%

of overall AI software spending is forecasted to be on off-the-shelf AI governance software, reaching $15.8 billion by 2030.[v]

Forrester

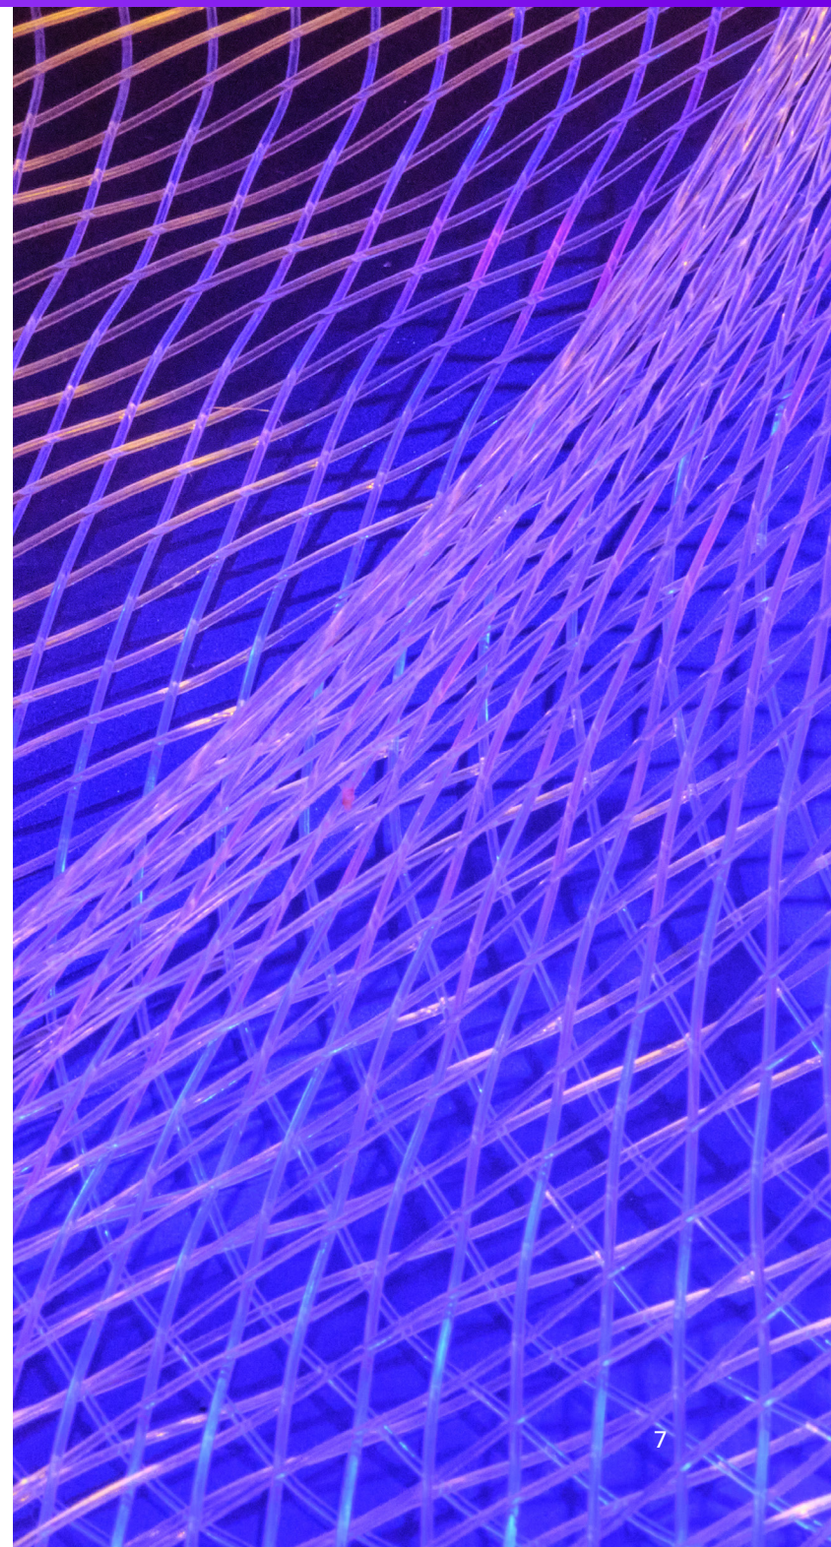

Key steps:

1. **Establish clear data guidelines** that lay out what data can and cannot be used (e.g., names, addresses, phone numbers, medical data, payment data, etc.) in generative AI models and implement quality controls to protect data privacy and intellectual property, as well as reduce the chance for bias and discrimination in results. For instance, organizations can dictate the use of anonymized or synthetic data for certain functions or require users to employ algorithms to identify and remove biases from data before training models.
2. **Institute a formal process to approve models** before deployment to ensure they meet business needs, align with the business architecture, and comply with security and regulatory requirements. Everyone across the organization should know what to look for when considering a model to be able to evaluate potential biases, ethical concerns, adherence to responsible AI principles, and risks. It should also provide transparency into who was involved in the decision and how that decision was made to create accountability.
3. **Manage model lifecycles.** The roles and responsibilities of each step need to be clearly defined and followed, including design, testing, development, deployment, ongoing monitoring and maintenance, and depreciation, to ensure the ongoing relevance and integrity of that model and reduce orphaned technology that can open the organization up to risks.
4. **Ensure accountability and continuous oversight.** A combination of technology and humans should be regularly monitoring and auditing results to verify AI-generated content and uncover potential problems before they can do harm. Responsibility for AI outcomes needs to be assigned, so it is clear who to go to when something goes wrong. The goal is to be able to clearly explain the decision-making processes of generative AI systems and provide transparency and confidence in outcomes.

> "Now is the time to start building a center of excellence for AI to enable people who aren't data scientists to be able to connect to AI and access these things in their products in a compliant way."

Laura Ellis,
VP, Data Engineering & Platform Analytics, Rapid7[vi]

RAPID7

ITU's *AI for Good Global Summit* brought leaders together to move from principles to practical AI governance frameworks.

Read more:
https://www.itu.int/hub/2024/07/key-findings-on-the-state-of-global-ai-governance/

4. Put in guardrails

Generative AI creates both great potential and risk. It can be targeted by attackers to give them entry into your organization, as well as used by bad actors to increase the effectiveness of their attacks. As a result, it is important to ensure your use of generative AI creates value, not unnecessary risk, for your business, and helps strengthen your defenses, not weaken them. This requires implementing guardrails that shut down potential attack vectors and promote the safe use of generative AI.

To start, you need visibility into all your generative AI usage – in your products and your operations – to understand your potential exposure. You need to understand the security built into your generative applications, as well as those you use or incorporate from partners/vendors to make sure all appropriate standards are being followed. To accomplish:

- Conduct a comprehensive inventory of all your generative AI usage to understand what data and models are being used.
- Categorize each application by purpose, complexity, risk level, and alignment with the established business approach to accurately assess the impact of each and monitor those that are high-risk or critical.

92%

of CISOs are reconsidering their future in the age of AI[vii]

Trellix

45%

of CISOs are developing an AI committee to review AI tools as well as implement governance, including security frameworks and standards.[vii]

Trellix

Next you need to implement a modern zero-trust framework with fine-grained controls to mitigate risks. Legacy solutions often use the urls for SaaS applications to make decisions, which covers most traffic and is very limited in what it reveals, forcing a binary decision (allow/deny) that is often too broad to be useful. What's needed is a solution that can identify and make decisions based on a variety of factors, including identity, device, location, application, instance, activity, behavior, and the data being used to enforce:

- Least privilege access
- Multi-factor authentication
- Encryption

Finally, you want to make sure you have additional security measures in place to help you protect against attacks at all levels – application, prompt, model, data. Some solutions to consider include:

- Data loss prevention (DLP) – prevents sensitive data from going into any of your LLMs. You also want to make sure any third-party AI applications you are using or integrating with meet your standards (e.g., privacy policy, encryption, protection).
- Enforcement tools – reminds users of appropriate usage (generative AI code of conduct) and flag potential violations, via a popup screen, to force users to actively acknowledge and accept responsibility for any potential risk they take.
- Detection and mitigation capabilities – allows you to identify potential malware, tampering, and other malicious activity and work to mitigate their impact. Capabilities you can deploy include behavioral analysis, anomaly detection, sandboxes, access controls (to prevent prompt or algorithm manipulation or unauthorized access to content), LLM firewalls, adversarial training, and testing protocols, among others.

Key steps:

1. **Understand usage** by cataloging what generative AI applications are being used, what data is being posted, and what the trends are within your organization.
2. **Implement a modern zero-trust framework,** using the necessary tools and policies to enforce appropriate usage and protect sensitive data to minimize liability and risk.
3. **Start red teaming** to identify where vulnerabilities exist and help prepare for real-time attacks.
4. **Update incident response plans** to include generative AI risks, so everyone knows exactly what to do in the event of a breach to minimize its impact and speed recovery.

Atlas

The MITRE adversarial threat landscape for artificial intelligence systems (ATLAS) outlines various attack methods that can be used to target AI, which include methods to evade models, perform a denial of service, erode the model integrity, conduct cost harvesting, and spam the AI system. The harm of these attacks can be financial, reputational, societal, user-related, or entail intellectual property theft.

Learn more:
https://atlas.mitre.org/matrices/ATLAS

5. Iterate, iterate, iterate

Start small to show teams how to use AI to do their jobs more effectively. This builds comfort and confidence in the technology and lays the foundation for bigger, more disruptive changes in the future. It's important to build a culture of experimentation and learning, so teams are not afraid to try and test new things. The goal is to be ready to fail and then quickly turn those failures into learnings that help you innovate and succeed.

> "Getting to a generative AI demo can be pretty easy, but the long tail of responses is what determines if it really works, so that's where you need to add more to your evals and start fine tuning to get to a point where you will consider putting it into production."
>
> **Naveen Rao,**
> VP of Generative AI, Databricks[viii]
>
> databricks

> "As we innovate, we let our customers and engineers try new things. Over the past year we released over 30 features across the entire platform. Each team has a small pod of people empowered to act almost autonomously against their objectives and key results (OKRs) and they tell us what their customers are saying. At the end of the day, it's about doing your research. To deliver a fantastic customer experience, you have to know what your customers want.
>
> We invest in our systems and internal tooling to ensure both new and iterative work can be done efficiently, enabling us to release code 5-10x per day for hotfixes, daily batch releases, and major weekly releases. Over the last week we changed 1,254 files with 224,000 lines of new code. We empower our engineers to change rapidly and make sure that if you're going to fail, you do it quickly."
>
> **Jamie Doheny,**
> General Manager, VP, Technology Services & Engineering, Braze[ix]
>
> braze

Generative AI gets better with feedback, using it to improve the accuracy, precision, and appropriateness of its results. This allows you to feel confident you know how it will respond in most situations in a production environment. The more you work with it, the more you understand and the more you can explore to uncover new lines of thinking or fine tune operations to better meet your needs and drive greater benefits.

Key steps:

1. **Start small** to build confidence and acceptance of new ways of doing things.
2. **Build a culture of experimentation** to make it okay to try and fail. Acknowledge any win (big or small) and highlight any learnings from failures to help build momentum and empower innovation.
3. **Iterate** to improve accuracy and refine outputs, as well as uncover new paths that can lead to even greater ideas and outcomes.
4. **Set metrics** that ensure you are measuring the effectiveness of everything you are doing and can gauge progress. This will enable unbiased evaluations of generative AI projects, helping teams determine what to pursue and what to potential abandon to drive desired results.

AWS helps you harness the potential of generative AI

With enterprise-grade security and privacy, access to leading FMs, and a broad array of services, tools and applications, AWS makes it easy to adopt, build, and scale generative AI to meet the needs of your use cases and customers. Please reach out for more information on how we can support you in your generative AI journey and create differential value.

Visit our website →
for more resources dedicated to software and technology businesses

Get in touch →
with our experts to help you begin or accelerate your generative AI journey

i. Source: World Economic Forum https://www.weforum.org/stories/2025/01/ai-2025-workplace/

ii. Source: Naveen Rao, VP of Generative AI, Databricks in roundtable titled, "Developing an AI Strategy for Software Companies" at the AWS for Software Companies Executive Forum, at re:Invent, December 2, 2024, Las Vegas, Nevada.

iii. Source: IBM Report https://www.ibm.com/thought-leadership/institute-business-value/en-us/report/business-trends-2025

iv. Source: Presented by Jonathan Shulman, Senior Partner, McKinsey and Company, in the keynote "Market Trends in the Software Industry," at the AWS for Software Companies Executive Forum at re:Invent, in Las Vegas, November 2024.

v. Source: Forrester https://www.forrester.com/report/global-commercial-ai-software-governance-market-forecast-2024-to-2030/RES181688

vi. Source: Laura Ellis, VP, Data Engineering & Platform Analytics, Rapid7 in panel discussion, "Everyone in My Organization Needs Data," during AWS for Software Companies: Generative AI and Data Executive Forum, in March 28, 2024, in New York, NY.

vii. Source: Trellix https://www.trellix.com/news/press-releases/92-of-cisos-question-the-future-of-their-role-amidst-growing-ai-pressures/

viii. Source: Naveen Rao, VP of Generative AI, Databricks, in roundtable titled "Developing an AI Strategy for Software Companies" at the AWS for Software Companies Executive Forum at re:Invent, December 2, 2024, Las Vegas, Nevada.

ix. Source: Jamie Doheny, General Manager, VP, Technology Services & Engineering, Braze, in a Lightning Talk titled " Delivering Exceptional Customer Experiences through Innovation", on May 15, 2024, during AWS for Software Companies: Generative AI and Business Applications Executive Forum, in New York, NY.