Modernize your software company with the disruptive power of generative AI

A helpful guide for business, product, and technical leaders

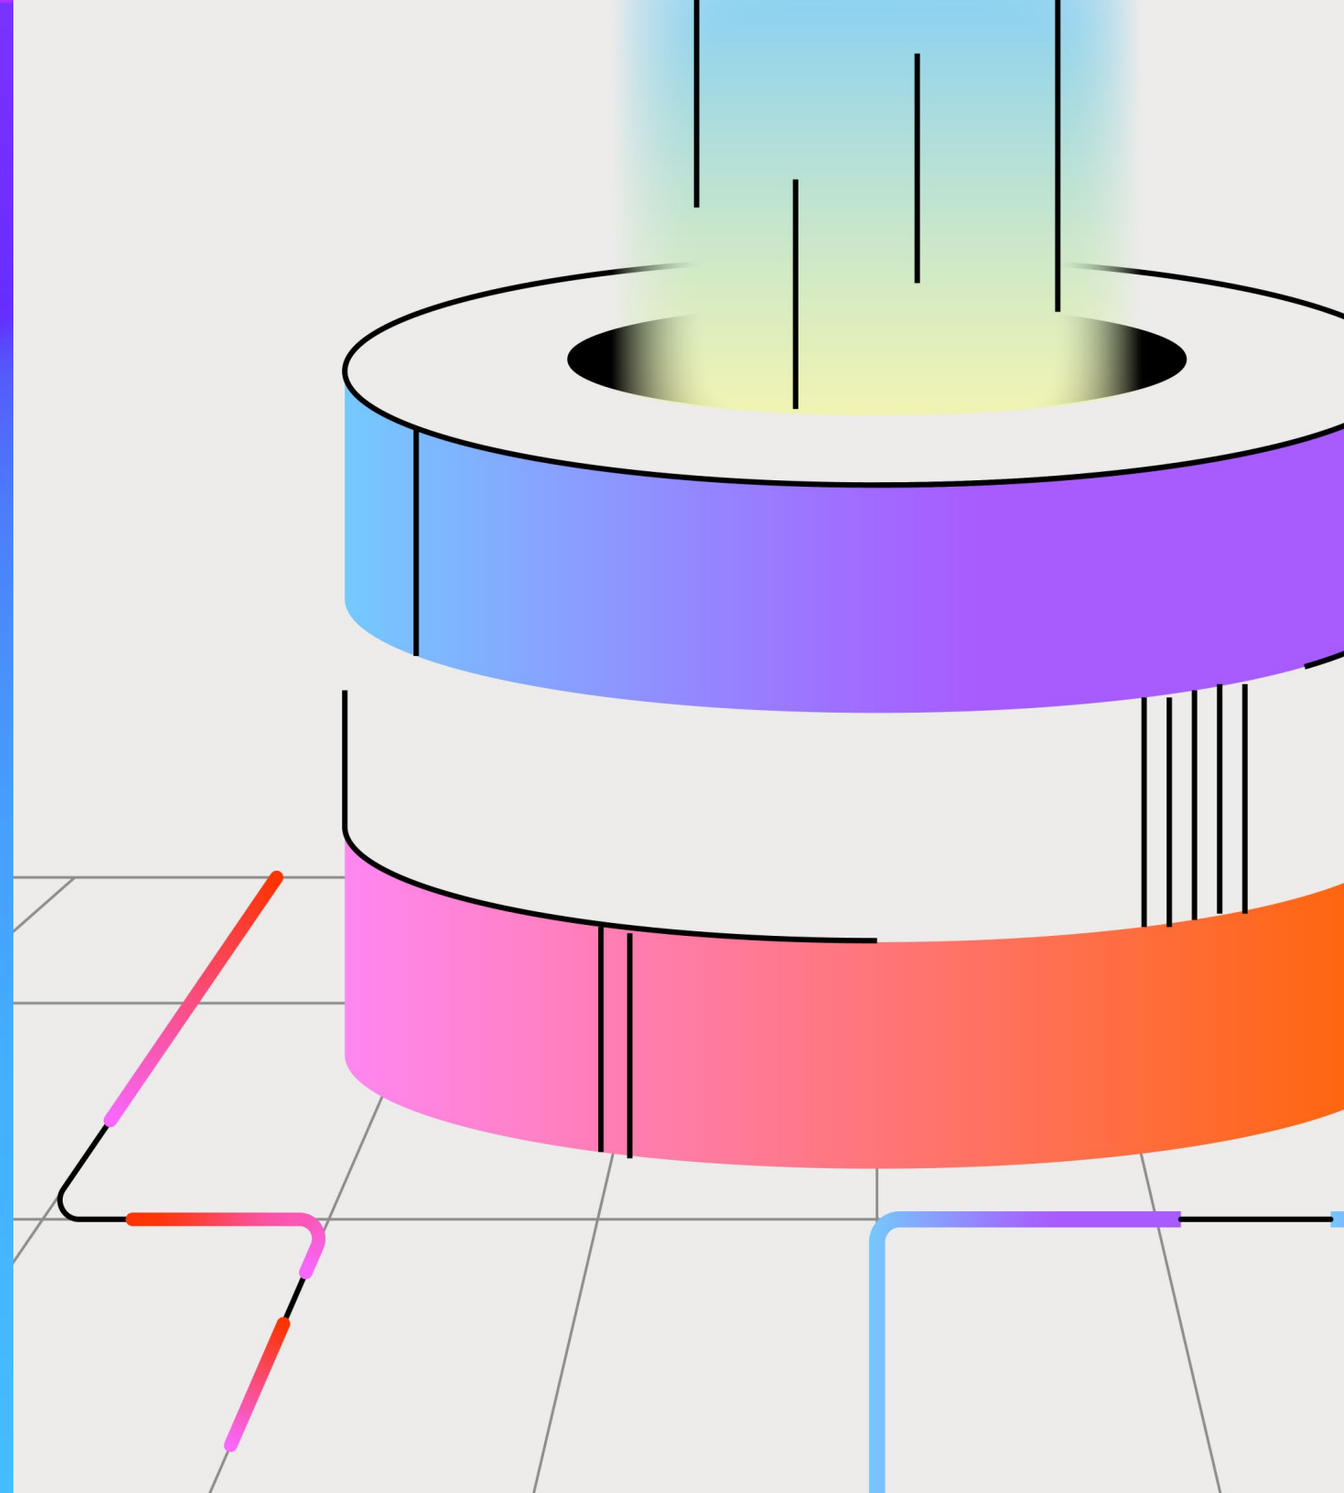

Table of contents

Succeeding in the era of generative AI

Generative artificial intelligence (gen AI) has captured the imagination of industries of all sizes, each looking to unlock its power and potential. By 2026, Gartner predicts that roughly 80 percent of enterprises will have deployed gen AI applications.[1] From unlocking new growth opportunities and reinventing customer experiences to boosting employee productivity, software and technology companies can use gen AI to drive impact and shape their future.

Despite the forward momentum and compelling opportunities, leaders of software and technology companies also see the challenges gen AI presents and have questions: Where does the technology fit within the business? How can we extract its value while minimizing risks? Do we have the right foundation to successfully execute on gen AI use cases? The good news is that software companies are already driving meaningful business value quickly with gen AI on Amazon Web Services (AWS).

In this eBook, you will find best practices and real-world examples to help you build the right strategy and strong foundation for gen AI success. From identifying the right use case to measuring results, this proven path can help your software company harness the power of gen AI to fuel innovation and gain a competitive advantage.

2 of 3 organizations are increasing gen AI investments due to strong early value to date[2]

"Almost all software categories are likely to have some impact from gen AI.... New use cases and features driven by gen AI are likely to spur revenue growth."[3]

[1] "Gartner Says More Than 80% of Enterprises Will Have Used Generative AI APIs or Deployed Generative AI-Enabled Applications by 2026," Gartner, October 2023
[2] "Now decides next: Moving from potential to performance," Deloitte's State of Generative AI in the Enterprise Quarter three report, August 2024
[3] Schneider, J., et al., "Navigating the generative AI disruption in software," McKinsey & Company, June 2024

Customer spotlight

Adobe successfully embeds generative AI into its product suite to better serve customers

About

Adobe is at the forefront of groundbreaking technology, offering a wide range of software products for digital content creation, management, and publication to bring any digital creative experience to life.

Challenge

Adobe set out to rethink its creative tools and embrace the power of gen AI to help its customers navigate this technological disruption. It was looking to integrate gen AI into its flagship creative software products and be in control of the quality and capabilities of its models by training its own Firefly family of gen AI models.

AWS solution

Using Amazon Elastic Compute Cloud (Amazon EC2) P5 instances, Amazon EC2 P4d instances powered by NVIDIA GPUs, Amazon Elastic Kubernetes Service (Amazon EKS), and Amazon Elastic Block Store (Amazon EBS), Adobe created an AI superhighway on AWS. Human-generated data was fed back into the training process, which helped AI models generate assets that resonated and were tailored to customer preferences. Amazon EC2 Reserved Instances (RI) helped Adobe scale up model training.

Results

Using AWS, Adobe was able to infuse gen AI into its product suite to improve its customer experience. Adobe achieved rapid innovation, launching the Firefly family of gen AI models in only nine months. To meet the spike in demand, Adobe scaled up training by 20 times in six months. By running its AI infrastructure on AWS, Adobe can dedicate its focus on differentiation.

Read the full story →

> "If you can't attract the AI research, engineering, and applied ML talent, look at AWS AI services like Amazon SageMaker and Amazon Bedrock. Focus on the differentiating generative AI value on top, as rebuilding all the foundational pieces will be extremely difficult."
>
> Alexandru Costin, VP, Generative AI & Sensei, Adobe

Benefits for software businesses

Gen AI helps software companies enhance existing applications and create new AI-driven products. It also drives spending on tools and infrastructure that support AI development—particularly those that boost productivity across IT and administrative teams, customer support and experience practices, and operational efficiency and process optimization efforts.

Benefits for IT and administrative teams:

- 86 percent more time spent focusing on innovation, such as automating server tasks[4]
- 31 percent less time spent on data entry, report generation, and other administrative tasks[5]
- Frees human agents to concentrate on resolving more complex customer issues and enhancing the overall customer experience
- Facilitates the creation of innovative, new products and services

Benefits for customer support and experience practices:

- Improves customer interactions and satisfaction by handling routine inquiries
- Analyzes consumer data to personalize customer recommendations and responses
- Enhances customer service with more sophisticated chatbots and virtual assistants

Benefits for operational efficiency and process optimization:

- 81 percent annual increase in the overall output of new applications developed[5]
- 134 percent more new features developed per year[5]
- Increases performance across business functions, automating routine creative tasks, accelerating ideation, and enhancing engagement

Beyond these larger business benefits, gen AI significantly boosts developer productivity, reduces development costs, and speeds up software processes, which can lead to rapid innovation and better adaptation to customer needs.[6]

[4] "The Business Value of Amazon Web Services," IDC, A Business Value Snapshot, sponsored by Amazon Web Services, June 2022
[5] "Discover the business value of cloud migration with AWS," AWS Infographic, 2024
[6] "Navigating the generative AI disruption in software," McKinsey & Company, June 2024

Achieving quick wins

As gen AI continues to evolve, software companies are uncovering fresh ways to harness its potential more quickly and efficiently than ever. Attaining these "quick wins" goes beyond understanding what gen AI *can* do; it requires effectively integrating it into business operations to drive meaningful outcomes. The path to success lies in the ability to establish quick wins that help your teams build experience and momentum for bigger and better successes.

Driving value with artificial intelligence

Software use cases that foster quick wins:

- **Employee productivity:** Start by implementing AI in internal tools, like coding and virtual assistants, where risk is minimized. Software coders can complete tasks up to twice as fast with gen AI.[7]
- **Process automation:** Use gen AI to automate key business processes that involve repetitive tasks, from customer support interactions to optimizing internal workflows.
- **Product development:** Gen AI enables the creation of innovative, new products that meet emerging customer demands for AI-enhanced features and drive additional revenue.
- **Marketing insights:** Enhance marketing efforts by generating personalized content and insights that allow marketing teams to efficiently reach the right audiences and improve engagement rates.

[7] "Unleashing developer productivity with generative AI," McKinsey & Company, June 2023

Barriers to successful adoption

Gen AI is expected to drive huge growth in the software industry, with spending potentially reaching up to US$250 billion by 2027.[8] However, this growth may cause significant disruption, primarily due to a higher rate of vendor switching. This increased turnover will be driven by new competitors using gen AI to challenge established companies, lower costs of switching software, and businesses' desire to stay current with the latest innovations.

Bypassing common obstacles

As gen AI continues to evolve, software companies must find simple, cost-efficient ways to harness its potential. The challenge is not merely in understanding what the technology can do but in effectively integrating it into existing business operations to drive meaningful outcomes.

Adding value with artificial intelligence

Reasons to choose AWS:

- **Greater choice:** AWS empowers organizations with choice and flexibility, offering a comprehensive set of AI capabilities from chips to foundation models (FMs) to applications, enabling rapid innovation.
- **Data as a differentiator:** With AWS, you can leverage your data to customize FMs, tapping into your existing data foundation on AWS, built-in governance controls, and customization techniques.
- **Built-in security:** AWS is architected to be a secure cloud computing environment for gen AI applications, offering robust tools to help protect each layer of the AI stack.
- **Unlimited scale:** Leverage decades of experience, a vast global infrastructure, and seamless integration across diverse data sources to deliver meaningful value to customers and employees at scale.

[8] "Navigating the generative AI disruption in software," McKinsey & Company, June 2024

6 steps for getting started with artificial intelligence

As a software or technology company navigating the dynamic landscape of AI, it's crucial to ask the right questions to better align your efforts with genuine business needs and deliver real value. Whether you're yet to migrate to the cloud or operating in a hybrid environment, here are the key questions to guide your gen AI journey:

Step 1: Create a culture of innovation and security → How are teams working together to unlock the potential of gen AI?

Executives at the highest level must take a wide-scale view to set goals, bring business and technical teams together, and commit to responsible AI from the onset. A people-centric approach can help educate a workforce on responsible AI and encourage experimentation while managing risks.

Step 2: Make data your differentiator → Do we have a data strategy in place to use organizational data for gen AI?

Data is your strategic asset—the difference between a generic gen AI application and one that truly knows your business. Effective data management, including privacy and security measures, is essential for successful AI implementation.

Step 3: Identify the right business problem to address → Do we have a specific business problem to solve that can demonstrate the value of gen AI?

There are several questions to address before embarking on a gen AI journey, including whether the technology solves a real business problem, if the project can benefit from AI, and if the project is a priority.

Step 4: Upskill teams → Are teams armed with the right skills to succeed in the era of gen AI?

Upskilling your workforce with gen AI skills can increase efficiency, creativity, and innovation. AWS offers practical skills that can be applied immediately to help maximize the abilities of existing staff.

Step 5: Scale beyond pilot projects → How will we sustainably scale AI across our organization?

Creating a center of excellence can help rally your community to push for new initiatives. AI can become an integral part of your yearly planning processes, bringing domain and technical experts together to brainstorm and determine next steps.

Step 6: Measuring the results and impact → How will we measure the success of our gen AI efforts?

Gen AI initiatives can be viewed through the lenses of agility, competitive advantage, and risk tolerance. AI efforts can be measured on what success looks like for your business, taking into account the processes being optimized.

Customer spotlight

Intuit delivers highly personalized customer experiences using the power of generative AI on AWS

About

Intuit provides financial management solutions to help over 100 million consumer and small business customers.

Challenge

Intuit wanted to tap into the potential of gen AI to enhance its financial products and services for small businesses and consumers. However, building and scaling gen AI applications required access to leading large language models (LLMs), robust data governance, and the ability to efficiently manage accuracy, latency, and cost.

AWS solution

Intuit built a proprietary gen AI operating system called GenOS on AWS, leveraging Amazon Bedrock and Amazon SageMaker. GenOS provides a unified platform to design, prototype, and deploy gen AI experiences using the broad library of LLMs in Amazon Bedrock alongside Intuit's financial data and custom LLMs. This allowed rapid development of personalized, accurate gen AI apps while maintaining data privacy and optimizing for low latency and cost-efficiency by hosting models on Amazon SageMaker.

Results

With Amazon Bedrock and Amazon SageMaker, Intuit could access leading LLMs, train custom models, and scale its gen AI infrastructure seamlessly. This powered the launch of Intuit Assist, a gen AI assistant embedded across products like TurboTax and QuickBooks. Now, Intuit can make 65 billion machine learning (ML) predictions per day and drive over 810 million AI-powered customer interactions.

Watch the customer story →

> "…we're delivering unparalleled experiences for our customers. Today, we've been able to achieve incredible scale with AWS running all of our data capabilities…"
>
> Nhung Ho, VP of AI, Intuit

2 paths for software companies

There are two pathways to gen AI for software companies: one for companies that have started their cloud journeys but may not yet have a significant portion of their workloads in the cloud and another for companies that have already started the migration process with a mix of on-premises and cloud assets.

For companies yet to migrate their workloads to the cloud:

Start with internal applications: This initial step allows for a controlled environment where security and transparency are prioritized. Start with internal applications, such as customer support chatbots, document generation, or internal code review tools; these can significantly benefit from the capabilities of gen AI. These internal applications are designed for employees only.

Focus on security: To adopt gen AI at scale while maintaining security, organizations should integrate security into their Governance, Risk, and Compliance (GRC) frameworks, conduct security assessments informed by cyber intelligence, and align security with best practices. Protect every layer of the AI stack, including the data, FMs, applications, and identity access controls, to create a robust and secure gen AI environment.

Transparency and testing: Transparency in AI decision making is crucial for building trust within the organization and with end users. Leverage services that offer tools for various stages of AI development, including model training, validation, and deployment, ensuring that AI applications perform reliably under various conditions. Extensive testing, including A/B testing and stress testing, can be conducted to identify potential issues and optimize AI performance.

For companies with a mix of on-premises and cloud assets:

Leverage pretrained models: Utilize pretrained models to accelerate gen AI deployment. Customize models to meet business needs and integrate with existing applications.

Implement scalable solutions: Use cloud resources to handle intensive AI workloads with scalability and flexibility. Maintain critical operations on premises for control and compliance.

Enhance data integration: Implement solutions for data integration between on-premises and cloud environments. Use data pipelines and extract, transform, and load (ETL) processes for enhanced data consistency and availability.

Focus on security and compliance: Establish security measures that are in place across both environments. Implement encryption, access controls, and regular audits to protect data integrity and comply with regulations.

Invest in training and upskilling: Equip your team with the skills needed to manage and optimize hybrid AI solutions. Encourage collaboration between cloud and on-premises teams for smoother operations.

Monitor and optimize: Regularly monitor the performance of gen AI applications across both environments. Use analytics and insights to continuously optimize operations and improve efficiency.

Customer spotlight

Omnicom collaborates with AWS to transform ad campaign development with generative AI

About

Omnicom Group is the world's leading global marketing and corporate communications company—with over 5,000 clients in more than 70 countries.

Challenge

Omnicom Group was looking to reimagine creativity and accelerate innovation with gen AI FMs for advertising. In particular, the organization wanted to accelerate the transformation of advertising campaign development.

AWS solution

Omnicom Group is advancing its AI-powered platform, Omni, to expand how it leverages its industry-leading datasets. It is adding capabilities powered by out-of-the-box FMs using AWS gen AI and ML services.

These include Amazon Bedrock and Amazon EC2 Trn1 instances powered by AWS Trainium chips, which are purpose-built for ML workloads.

Results

The Omni platform uses AWS to ingest trillions of advertising signals monthly and store more than 10 petabytes of data. It helps automate activities such as developing creative briefs, building media plans, segmenting audiences, and measuring advertising performance.

Read the article →

> "The innovation, security, simplicity, and reliability of AWS machine-learning and generative AI services help us accelerate our data capabilities to deliver transformational innovation and results to our customers."
>
> Jonathan Nelson, CEO, Omnicom Digital, Omnicom Group

How AWS can help

AWS delivers the most comprehensive set of capabilities at every layer of the stack to help you innovate with gen AI, including the most performant and cost-effective infrastructure for training and running ML models at scale, the easiest way to build gen AI applications with LLMs and other FMs, and the most capable gen AI–powered assistant to help transform how work gets done—and all with enterprise-grade security and privacy built in.

AMAZON Q

Transform how work gets done with generative AI

Amazon Q is the most capable gen AI–powered assistant for accelerating software development and leveraging companies' internal data to increase productivity, creativity, and efficiency: Amazon Q Business is a gen AI–powered assistant that lets users ask complex questions, get comprehensive answers, and execute actions—all based on your enterprise's content, data, and systems. Amazon Q business easily connects to more than 40 popular platforms, like Microsoft 365, Salesforce, and ServiceNow, providing users with access to information securely according to their permissions. Within Amazon Q Business, employees can use Amazon Q Apps to quickly create and share lightweight AI apps instantly with natural language prompts—turning their ideas into apps based on your enterprise data.

Amazon Q Developer is a gen AI–powered assistant that helps developers and IT professionals with all of their tasks—from coding, testing, and upgrading to troubleshooting, performing security scanning and fixes, optimizing AWS resources, and creating data engineering pipelines.

At Amazon, Amazon Q Developer success means:[9]

79%

of AI-generated code reviews were shipped without additional changes

4,500

developer-years of work have been saved

$260M

in annualized efficiency gains using upgrades made by Amazon Q

[9] "Amazon CEO Andy Jassy Says Company's AI Assistant Has Saved $260M And 4.5K Developer-Years Of Work: 'It's Been A Game Changer For Us'," Yahoo Finance, August 2024

AMAZON BEDROCK

The easiest way to build and scale generative AI applications with foundation models

Amazon Bedrock is designed to integrate seamlessly with on-premises, cloud, and hybrid infrastructures. As a fully managed service, it provides a single API for accessing and utilizing high-performing FMs from leading AI companies including AI21 Labs, Anthropic, Cohere, Meta, Mistral AI, Stability AI, and Amazon. It offers a broad set of capabilities to build secure, private, and responsible gen AI applications. With Amazon Bedrock, you can:

- Experiment and evaluate the best FMs for your specific use case
- Easily access and leverage the latest gen AI innovations from leading AI companies
- Seamlessly switch between different FMs and upgrade to the latest model innovation with minimal code changes
- Customize FMs with your business' proprietary data using techniques like fine-tuning and retrieval augmented generation (RAG)
- Build agents that execute tasks using your enterprise systems and data sources
- Eliminate the need for infrastructure management using serverless architecture

Benefits of Amazon Bedrock:

Integrates seamlessly

with on-premises, cloud, or hybrid infrastructures

Enables you to get to market faster

with a fully managed service, without the need to build costly AI infrastructure or hire new talent

Offers a single API,

allowing you to focus on a single API, along with a broad set of capabilities you need to build gen AI applications easily, allowing your teams to focus on product innovation

AMAZON SAGEMAKER

Streamline your entire foundation model development process

Amazon SageMaker offers a comprehensive suite of tools for building and training ML models, including FMs, with billions of parameters. This powerful AWS service includes pre-built algorithms, frameworks, and a variety of tools designed to streamline the entire FM development process by:

- Providing all the modern tools needed to build, train, and deploy AI models, making it a comprehensive solution for FM development
- Supporting a wide range of algorithms and frameworks, giving software application teams the flexibility to choose the best AI tools for their specific needs
- Making it easy to scale AI applications as your business grows, ensuring that your solutions can handle increasing demands
- Offering purpose-built infrastructure for distributed training at scale, reducing FM training time by 40 percent

Amazon SageMaker is ideal for:

Software companies yet to migrate to the cloud

for a flexible and comprehensive toolset that automates the creation and customization of AI solutions from scratch

Software companies with hybrid environments

for maintaining the consistency and scalability of the entire AI lifecycle across on-premises and cloud infrastructure

Integrate generative AI to drive innovation across your software business

Software and technology companies are using gen AI to shape the future of business and work, accessing its tangible value to gain competitive advantage. AWS offers the foundation on which to build your gen AI initiatives, providing flexible and scalable AI services designed to be secure. Boldly innovate, drive growth, and deliver exceptional customer experiences while maintaining the highest levels of trust, privacy, and security. Learn more on your own or contact our gen AI experts and start your own AWS Cloud journey today.

Request a free consultation →