Generative AI for software companies: 6 steps to success

A proven approach to leveraging the full power of generative AI

Table of contents

INTRODUCTION

The opportunity for software companies in the era of artificial intelligence

When deployed with the right strategy, generative artificial intelligence (gen AI) has the potential to change the anatomy of work. According to a study conducted by McKinsey Digital, about 75 percent of the business value that gen AI will drive falls into four domains: customer operations, marketing and sales, software engineering, and research and development (R&D).[1] A few use cases stand out and are likely to deliver a significant return on investment (ROI) for software companies by helping to improve productivity down to the individual worker.

For example, developers can get code suggestions to improve productivity and develop applications faster. Marketing teams can create ad copy and stunning visuals in minutes using natural language instructions. Customer service departments can transform call centers using gen AI–powered chatbots to answer customer questions with conversational interfaces. Finance teams can produce reports tailored to different business units, making it easier to build and consume insights across organizations.

The lack of a single, universally accepted playbook for gen AI success is keeping some businesses on the sidelines, with otherwise decisive decision makers unsure of how to take the next (or first) step of the AI journey.

By democratizing gen AI, Amazon Web Services (AWS) helps software businesses overcome barriers to adoption, providing the expertise and services it takes to forge ahead with confidence. This eBook outlines a modern path—from taking the first step to measuring the results—with best practice insights from Amazon and its experience helping thousands of AWS customers realize their own AI initiatives.

In a global survey of software and digital native companies, all respondents said there is a need to increase investment in AI and, especially, to stay ahead of the competition. In fact, 83% agree that in the next 5 years, the only competitively viable publishers of business and consumer software will be those executing an effective AI product strategy.[2]

[1] "The economic potential of generative AI: The next productivity frontier," McKinsey Digital, June 2023
[2] "The Real State of AI/ML Capabilities in Software and Digital Native Companies," SoftServe, February 2023

What are artificial intelligence, generative AI, and machine learning?

AI is a way to describe any system that can replicate tasks that previously required human intelligence. Almost all AI systems are built using machine learning (ML).

ML utilizes large amounts of data to create and validate decision logic, forming the basis of an AI model. The AI application then feeds input data into that model, and the model outputs human-like decisions. The rapid advancement of ML, a massive proliferation of data, and easy availability of scalable compute capacity can help businesses accelerate how they use AI and, more specifically, how they leverage gen AI.

Gen AI is a type of AI that can create new content and ideas, including conversations, stories, images, videos, and music. Like most AI, gen AI is powered by ML models—large language models (LLMs) that are pretrained on vast amounts of data and commonly referred to as foundation models (FMs).

Why generative AI (and why now)?

Before diving into the six steps of a gen AI journey, let's explore why software companies should embark on that journey in the first place. According to Gartner, gen AI is the number one type of AI solution deployed in organizations, with 29 percent of the surveyed respondents reporting that they have deployed and are using gen AI.[3]

Globally, we have reached an inflection point at which most customer experiences and applications can be credibly reinvented with gen AI. Thus, the development of an AI strategy that includes gen AI is fundamental to achieving successful business outcomes. Even after completing the steps outlined in this eBook, AI teams will need to regularly remind themselves of what they're working toward—staying focused on the precise business benefits that can be unlocked by fully leveraging gen AI technology.

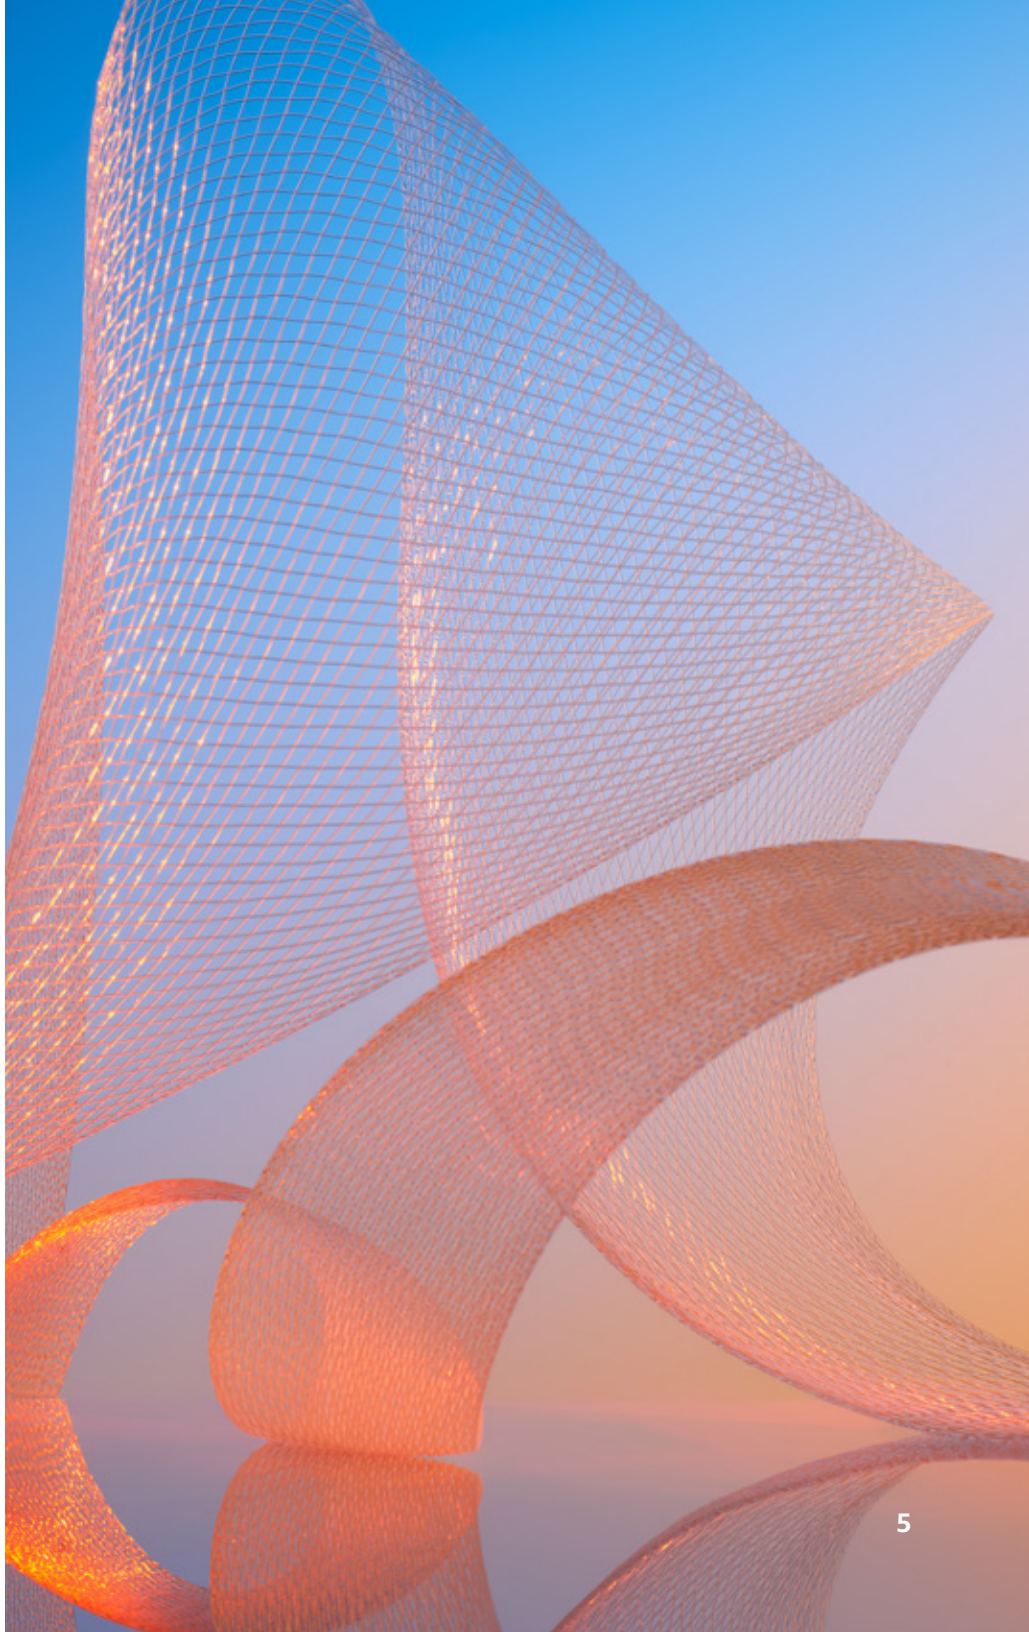

[3] "Gartner Survey Finds Generative AI Is Now the Most Frequently Deployed AI Solution in Organizations," Gartner, May 2024

Businesses worldwide are already realizing the impact of:

Boosting employee productivity

Gen AI is being leveraged for its transformative value to help software companies reach new levels of productivity for their customers. Advancements in the technology can be used to increase employee productivity with the help of AI-powered conversational search, content creation, text summarization, and code generation. Real-time monitoring company Dataminr used **AWS Inferentia** to grow its AI models while increasing processing speed, reducing costs, and improving accuracy—all of which led to bigger, better innovation on new projects and allowed the company to deliver faster, more accurate services.

Enhancing customer experiences

Today's organizations can take advantage of software solutions that incorporate gen AI to enhance customer engagement, increase personalization, and attract new users through deeper experiences. AWS customers have improved their own customer experiences through the effective use of chatbots, virtual assistants, intelligent contact centers, personalization, and content moderation. One of those customers—restaurant management platform **Upserve**—worked closely with the **Amazon Machine Learning Solutions Lab** to build a predictive model ML solution that was designed to predict how many people would visit per night and which menu items would be popular. All of this was made possible by integrating table management, point of sale, and other data systems in real time.[4]

Driving process optimization

Legacy business practices and other inefficient processes are critical areas for gen AI–driven transformation. Already, AWS customers are improving their business operations by leveraging intelligent document processing (IDP) to automatically extract and summarize data from documents and insights through AI-powered Q&A experiences. In the area of data augmentation, AWS is helping a growing number of organizations generate synthetic data for the training of ML models whenever the original dataset is small, imbalanced, or sensitive. And AWS customers across a range of industries are improving logistics and reducing costs by evaluating and optimizing different supply chain scenarios.

[4] "Upserve Case Study," AWS Case Study, 2016

The generative AI journey

is not necessarily a straightforward path. Achieving success with gen AI requires more than great technology—it also means ensuring that the organization is aligned to the right goals. Identifying and reaching those goals may necessitate broad changes in processes, management, and culture. In the next sections, we will explore how organizations can overcome challenges that often impede progress and take the right steps to implement AI in efficient, sustainable ways.

HOW IT WORKS

6 steps of a successful generative AI journey

Step 1: Champion an innovative culture

Unlocking the full business potential of gen AI requires cultural changes in team organization, objectives, and outlook.

For gen AI to proliferate through an organization, both business and technical teams have to work together and share the same priorities. To achieve this at the outset, the gen AI effort must be supported from the highest levels, with goals set by executive champions and an investment in the technology and processes that enable success. This includes, among others, a commitment to build more responsible AI from the beginning—working to identify and mitigate bias, improve explainability, and help keep data private and secure. By taking a people-centric approach, organizations can work to educate their workforce on responsible AI and build diverse teams to bring more perspectives to the table and improve fairness.

It's important for management to take a wide-scale view while fostering AI initiatives. Executives must be firm in their goals but also flexible in how the organization reaches them. Mistakes are sure to be made. But by staying focused on the long-term outlook and not allowing discouragement, organizations can glean wisdom from every error and apply those learnings to promote an innovative culture throughout the business.

Perhaps the largest cultural change that organizations must undergo is utilizing the opportunity that is inherent in mistakes. Integrating AI is an iterative process that can only succeed through constant experimentation. Often, these experiments will result in failure. Only by learning from mistakes—and refusing to grind progress to a halt in the name of determining "what went wrong"—can organizations consistently reach the breakthrough successes waiting on the other side.

How Amazon did it

Amazon has been using AI for more than 20 years. After a decade of leveraging the technology, our leadership team asked every business leader in the organization—irrespective of whether they ran a research team, a fulfillment center, or an HR organization—to answer the question of how they planned to use AI in their businesses. "We don't plan to" wasn't an acceptable answer in most cases, which forced the leadership, domain experts, and technical experts to collaborate on AI initiatives and let nothing halt their progress—even in instances where tangible benefits were still years down the road.

In addition to hiring external data scientists, Amazon created the Machine Learning University (MLU), which trained many of its developers to use AI more effectively. The company also built tools like **Amazon SageMaker**, which simplifies model creation and removes barriers to entry, enabling AI technologies and initiatives to scale more effectively. Additionally, Amazon created a set of pre-built **AI services** that provide ready-made intelligence to address common business use cases—without customers having to build their own models. For example, **Amazon Bedrock** is a service that makes gen AI FMs from AI21 Labs, Anthropic, Cohere, Meta, Stability AI, and Amazon accessible via an API. Amazon Bedrock is one of the easiest ways for customers to build and scale gen AI–based applications using FMs. Amazon Bedrock offers the ability to access a range of powerful FMs for text and images—including **Amazon Titan** FMs—through a scalable, reliable, and AWS managed service designed to be secure.

90%

of the companies listed on the Forbes 2024 AI 50 list are AWS customers.[5]

This is why major organizations—including Intuit, Thomson Reuters, AstraZeneca, Ferrari, Bundesliga, 3M, and BMW—and thousands of startups and government agencies around the world are transforming their operations, industries, and missions with gen AI solutions from AWS. We work to take these technologies out of the realm of research and experiments by helping to extend their availability far beyond a handful of startups and large, well-funded tech companies.

Let's take a look at some examples of how Amazon is leveraging AI.

For starters, Amazon uses AI throughout its fulfillment process and leverages a forecast system that can predict demand for nearly every product in its enormous inventory. These prediction models allow Amazon to better deliver on customer expectations of convenience, cost, and delivery speed.

"We forecast millions of products every single day across all of our Amazon sites worldwide," says Jenny Freshwater, director of forecasting at Amazon. "And without machine learning, we would not be able to produce those forecasts."

Amazon also created **Amazon Alexa**, which provides customers with an entirely new way to interact with technology. And the company developed groundbreaking technology with autonomous flight via Amazon Prime Air drones. Amazon also uses robotics in its fulfillment centers to get packages to customers faster. The examples go on and on.

Achieving these successes required significant investments in technology, research, and talent. But those investments would have gone to waste without the cultural changes that pushed them forward—despite many failures and unexpected challenges. Every organization must follow suit and foster this same fault-tolerant culture of experimentation and innovation before the AI journey can truly begin.

[5] "Forbes 2024 AI 50 List," Forbes, April 2024

STEP 2

Make data your differentiator

Data is a critical piece of gen AI. It's true that some general LLMs can be used "out of the box" for select use cases. However, when you want to build gen AI applications that are unique to your business needs, your organization's data is your strategic asset. Data is the difference between generic gen AI applications and those that truly know your business and customers and can deliver better, more differentiated experiences.

While some software companies will build and train their own LLMs with extensive amounts of data, many more will use their organizational data to fine-tune FMs for their unique business needs or to add context to prompts through retrieval augmented generation (RAG). For example, you can use data from sources such as your data lake, database, and data warehouse to create a chatbot that provides technical support for your products or a model that provides marketing copy trained on your best-performing ads. Because of this, success with gen AI requires relevant, high-quality data, which means that you need a strong data strategy in the cloud.

According to McKinsey Digital, "...companies that have not yet found ways to effectively harmonize and provide ready access to their data will be unable to fine-tune gen AI to unlock more of its potentially transformative uses."[6]

The right data strategy for gen AI includes a comprehensive set of services to store and query data at scale, break down silos so you have ready access to all of your data for gen AI applications, and make sure your data is secured and governed throughout the AI lifecycle. Keep reading to discover how a data foundation built on AWS gives you a strategic advantage.

"Our team has built a platform that trains, tests, and tunes machine learning models for companies in a variety of industries. For example, one of our previous projects was working with an online marketplace to automatically shorten product titles—which can often be 40–50 words long—so they could then be read aloud for a conversational shopping experience."

Joan Xiao, Lead Machine Learning Scientist, Figure Eight

[6] "What every CEO should know about generative AI," McKinsey Digital, May 2023

How Figure Eight did it

For over a decade, Figure Eight has helped product, data science, and ML teams more easily sort through their data for ML projects. Using its enterprise-ready platform, users can quickly prepare and categorize their own data—while Figure Eight annotates, judges, and labels that data to create a base for project models.

With AWS, Figure Eight has generated more than 10 billion data labels for real-world AI applications so far—all of which allow teams to focus on improving ML models' accuracy instead of managing, cleaning, and disputing data. Its platform provides accurate training data for likely scenarios by testing labelers before and during data tasks. This activity helps financial institutions keep their transaction models up to date, prevent data drift, and better train optical character recognition (OCR) models on financial documents.

STEP 3

Find the right business problem to solve

One common mistake software companies and others make in their AI journeys is employing discrete data scientists who work in silos to build models as proofs of concept rather than solving real business problems. This way of working has shown that many software and digital native companies are missing genuine business value—and that AI has so much more to offer.

Here are some important questions to ask before embarking on your AI journey:

- Is the project important enough to get attention and adoption?
- Does it solve a real business problem?
- Do we have the right data to solve the problem?
- Will the project benefit from AI?
- Can it eventually be operationalized?

For software company founders and their teams, the questions don't stop there. Many turn to the **AWS Resources for Software Companies** page for access to technical and business insights in the form of software industry case studies, timely and relevant podcast episodes, workshop materials, on-demand webinars, and more.

In a successful AI journey, organizations create AI teams to address specific business problems. This requires including both technical and domain experts within these teams. While the technical experts will take on the brunt of model creation, they need the field knowledge of domain experts to define precise business challenges and identify the data most important to finding a solution.

This approach is also critical to change management. When technical and domain experts collaborate to create models, employees will feel more confident in making decisions based on the algorithm's logic.

Together, these teams should also work through how to measure success. Freshwater says the companies should make sure that they have "very crisp and clear metrics as you embark on the machine learning journey. Many times, your models are taking over for something existing, and you want to make sure that they're actually better and that you can measure it."

For more on measuring the success of AI initiatives, refer to Step 6 in this eBook.

Some organizations have the talent in-house to identify AI-addressable problems and implement the appropriate pilot programs. Software organizations can leverage the **AWS Generative AI Innovation Center** to work backwards from business challenges and follow each of the steps for creating AI projects and initiatives.

How Tangent Works did it

Tangent Works specializes in automating ML modeling processes and leverages AWS to provide accessible, cost-effective AI technology for businesses without dedicated data science teams.

By using AWS, Tangent Works efficiently launches services, manages computer and storage needs, and supports rapid growth. Its Tangent Information Modeler (TIM) empowers business users and data scientists to easily create accurate AI models for time-series data. AWS then enables Tangent Works to scale resources dynamically to meet different customer demands, while tools like **Amazon Elastic Kubernetes Service** (Amazon EKS), **AWS Fargate**, and **Amazon Relational Database Service (Amazon RDS) for PostgreSQL** can then concentrate on facilitating workload management.

TIM offers bespoke AI capabilities, delivering rapid, cost-effective model creation and execution. This technology is widely applied across multiple industries for tasks like demand forecasting, maintenance scheduling, and fraud detection. The simplicity of updating and running models allows customers to keep insights current and adapt quickly to changing conditions.

Working with Siemens Digital Industries Software, Tangent Works integrated its technology into Siemens' industrial Internet of Things (IoT) platform, enabling users to analyze IoT data and develop data models for improved decision making. Tangent Works also aims to enhance automated modeling functions within existing AWS ML tools and refine its time-series anomaly detection technology.

By utilizing AWS, Tangent Works facilitates widespread access to advanced analytics and ML capabilities so organizations can harness insights without the need for extensive in-house data expertise.

STEP 4

Upskill your teams

In parallel with creating a comprehensive data strategy, today's software companies must focus on arming their teams with the right skills to succeed in the era of gen AI. However, businesses across dozens of industries are growing aware of an expanding skills gap—the separation between new technologies and the ability of internal IT and business specialists to take full advantage of them.

It's an issue that should set off alarm bells, with many thought leaders calling this monumental shift "the reskilling revolution." In response, nearly half (47 percent) of 37 organizations with gen AI tools offer formal training to their IT teams, while 38 percent are planning to implement such training.[7] Additionally, while hard skills in AI will continue to be important, soft skills—like creativity, ethics, and adaptability—will grow increasingly critical with gen AI.[8]

To help individuals train for the AI jobs of the future, AWS released on-demand skills training to support those who want to understand, implement, and begin using **gen AI**. Amazon has also designed training courses specifically for developers who want to use **Amazon Q Developer**.

While there's no one-size-fits-all solution to the AI skills gap, there are proven methods that can maximize the abilities of existing staff, reducing the need to make large investments in recruiting pretrained expert talent.

[7] "Generative AI Training for IT Organizations," Gartner Peer Community Report, Dec 2023–Jan 2024
[8] "The future of jobs in the age of AI, sustainability and deglobalization," World Economic Forum, May 2023

These methods include:

- **Defining the skills gap:** Before closing the skills gap, an organization must identify the precise differences between what it needs or wants its employees to do and what its employees currently can do.
- **Understanding how skills are mapped:** Because AI initiatives are interdisciplinary efforts, an organization should map the skills needed across data scientists, data engineers, business analysts, application developers, statisticians, and other subject matter experts in the business.
- **Customizing training for specific needs:** If an organization has existing training curriculums that could be useful, it should work to tailor those materials to the business' specific AI needs. Leaders should also investigate pretrained AI services that provide ready-made intelligence for business applications and workflows.

In addition to training, organizations need to align teams to successfully tackle AI problems. They can achieve this by:

- **Promoting a culture of empowered teams:** AI project teams must be cross-functional, with the authority to execute individual objectives and the freedom to organically cross-pollinate with other teams as demands dictate and opportunities arise. To make this kind of teamwork possible, management will need to embrace new structures—letting go of the strictly hierarchical and departmentally siloed organizational models of the past.
- **Starting with a pilot team:** The business should establish a pilot team of engineers, IT and AI practitioners, and line-of-business leaders—and task it with an AI project. "I'd recommend putting a couple of really smart people on trying to figure out what metrics you want to optimize for or predict," says Freshwater.
- **Enabling organic transformation:** Once the pilot project is complete, the business can split up the team, add people to create new teams, and task them with new projects. This process continues, allowing knowledge to spread organically from veteran team members to new recruits and pollinate between teams.

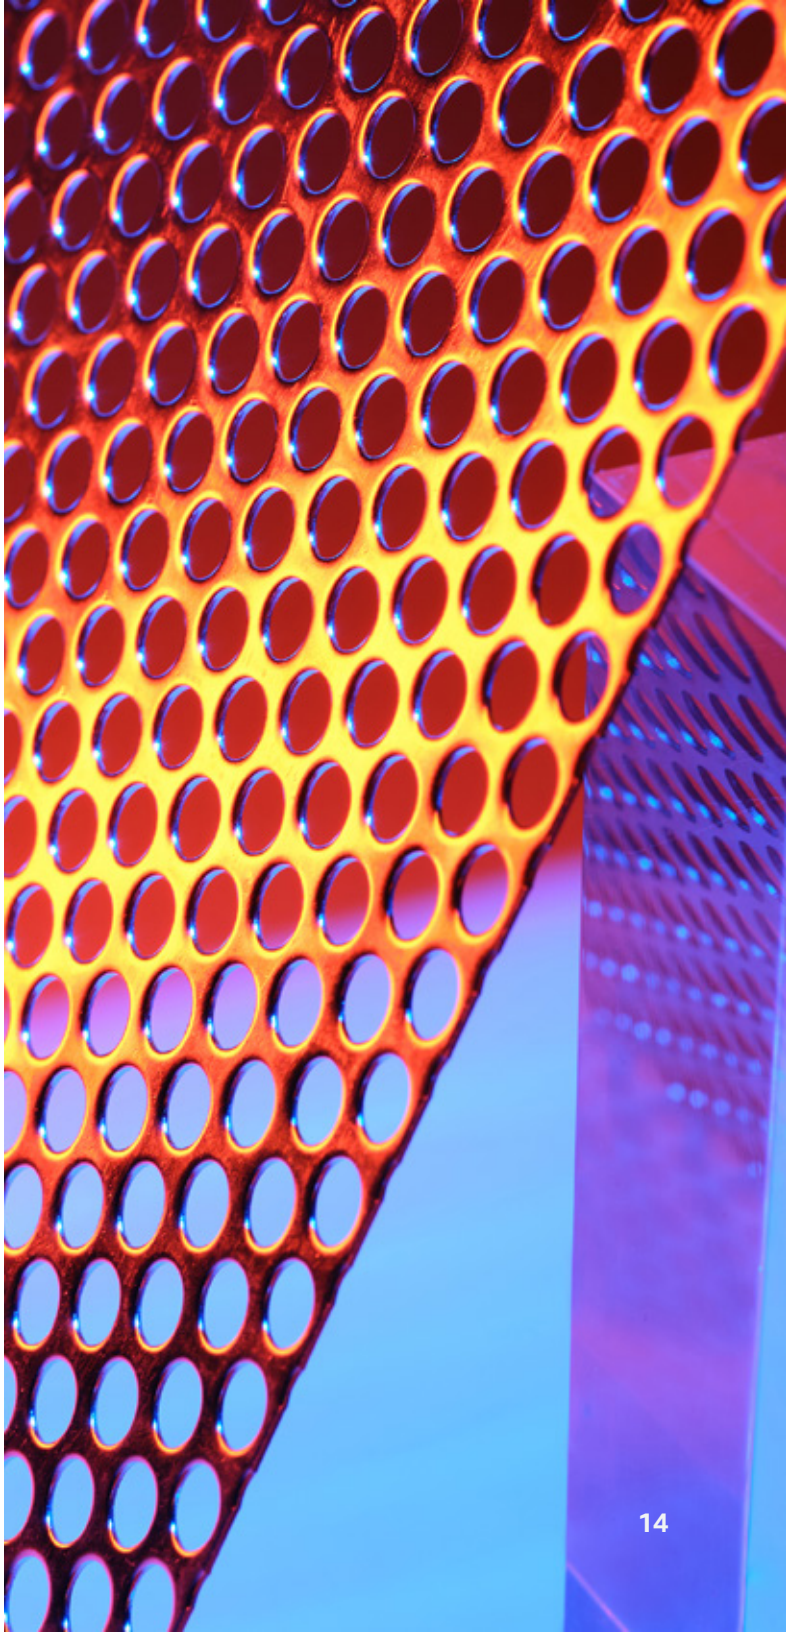

How Smartsheet did it

Smartsheet is the modern enterprise work management platform trusted by millions of people at companies across the globe, including approximately 85 percent of the 2023 Fortune 500 companies. The company was seeking ways to leverage AI to streamline knowledge management and accelerate employee productivity in the cloud.

Amazon Q is embedded in the company's Slack app to help thousands of global employees quickly get answers to their questions without having to know where the information lives organizationally. The Smartsheet Amazon Q chatbot, @AskMe, is trained on public help documents, training courses, and hundreds of all-employee Slack help channels. Employees can simply tag @AskMe in any Slack channel and ask a question, and Amazon Q will instantly give them an answer.

"Smartsheet and AWS share a commitment to innovation that simplifies work," says Rich Geraffo, vice president and managing director at AWS. "With Amazon Q, Smartsheet is applying generative AI to make its organizational knowledge more accessible, boosting employee productivity and freeing up its team's time to focus on delivering for customers. We look forward to continuing to work with Smartsheet on new AI-powered solutions that empower employees and customers to achieve more."

"Previously, it was too difficult for our 3,300 employees to find the information they needed across public help documents, training courses, and hundreds of all-employee Slack help channels. Now, employees can simply tag @AskMe in any Slack channel, ask a question, and Amazon Q will instantly give them an answer."

Bani Bedi, SVP of Corporate Development & Strategy, Smartsheet

STEP 5

Scale beyond pilot projects

After the first few successful pilots, AI teams must take the next step on the journey: sustainably scaling AI across the business. This is both a technical and a cultural challenge. There are several ways software companies can approach the cultural shift necessary to scale AI. Some might find success by creating a center of excellence that rallies the community and continues to push for new initiatives. Or, like Amazon, organizations can make AI an integral part of yearly planning processes, continuously bringing domain and technical experts together to brainstorm and plan next steps.

Achieving scalability requires organizations to help developers use AI. Building models at scale can be labor-intensive, which can slow innovation. With the Amazon Bedrock serverless experience, you can experiment with and evaluate top FMs for your use case, privately customize them with techniques such as fine-tuning and RAG, and build agents that execute tasks using your enterprise systems and data sources. Amazon Bedrock helps you rapidly adapt and take advantage of the latest gen AI innovations with easy access to a choice of high-performing FMs like **Amazon Titan**. Regardless of the models you choose, the single-API access of Amazon Bedrock offers the flexibility to use different FMs and upgrade to the latest model versions with minimal code changes.

Many organizations are scaling through **AWS AI services**, a set of pretrained and managed services that can be used as building blocks to address common use cases, including personalizing recommendations, modernizing contact centers, improving safety and security, and increasing customer engagement.

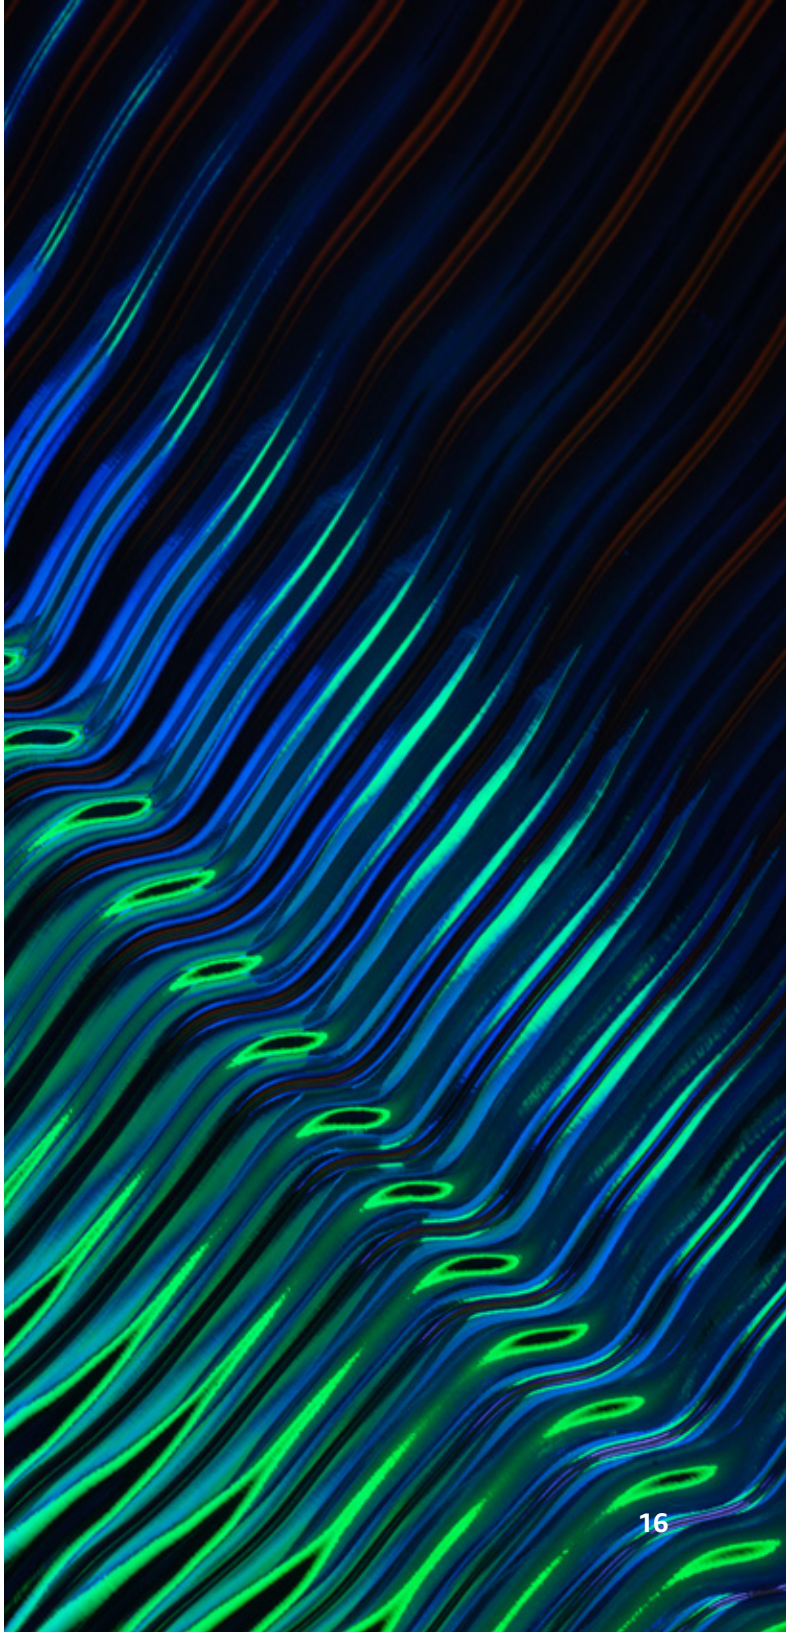

How Showpad did it

Showpad is a global leader in sales and revenue-enablement technology. It brings together industry-leading training and coaching software with innovative content solutions to empower sales and marketing teams to engage buyers.

To meet new expectations, Showpad needed to evolve its seller roles. This would help prospects and customers make sense of new information and put it into context for their businesses and the specific challenges they face. With a bit of AI innovation throughout the product experience, Showpad introduced **Amazon Bedrock** and **Amazon SageMaker** to test the waters and subsequently push new models to production.

What came out of this process was PitchAI, Showpad's latest sales coaching software.

The software helps sellers improve their sales pitches with actionable feedback in real time. Its AI-powered search enriches search results and reduces sellers' time spent searching for the information, while its AI-powered test questions streamline the process of creating training materials for sellers.

"Showpad is leveraging AI to improve how sales and marketing teams partner together in pursuit of a shared goal—closing more deals. It enables sellers to focus more time on what truly matters—bringing meaningful value to each buyer interaction."

Tony Grout, CPO, Showpad

STEP 6

Measure the results

When measuring the results of AI efforts, the traditional "project ROI" viewpoint—where a project has defined start and end points, a budget, and a return—is reductive and can be detrimental to the initiative's success. If the project doesn't generate a positive return within the given time frame, the business may lose interest and miss out on critical opportunities down the line. Instead, both business and IT decision makers must measure AI efforts based on what success means for their businesses with regard to the processes being optimized. In addition, they must view AI efforts as long-term investments, acknowledging that any true "return" may not be realized for several years and not before countless iterations.

When planning AI initiatives, it's better to view the process through the lenses of agility, competitive advantage, or risk tolerance rather than "expected" return. An organization will have greater success if it disregards the question of "What will be my return on investment in X months?" in favor of something more like "If we don't invest in this now, will we fall behind our competitors in X years when the technology matures?"

While traditional ROI metrics may not be the best approach, the business impact of AI initiatives can still be measured—it simply requires a different outlook. AI results can be measured through something resembling a value tree, where the main trunk of the tree represents the traditional revenue return, and branches extending from the trunk recognize the value of other business outcomes. The specific branches of the value tree will depend on the organization, the industry, and the initiative. But they might be things like time saved through automated processes; new leads, markets, and opportunities identified; customer service improvements; or increases in upsells.

Measuring the success of AI through a more holistic and long-term model will keep organizations focused on the best outcomes for their business future.

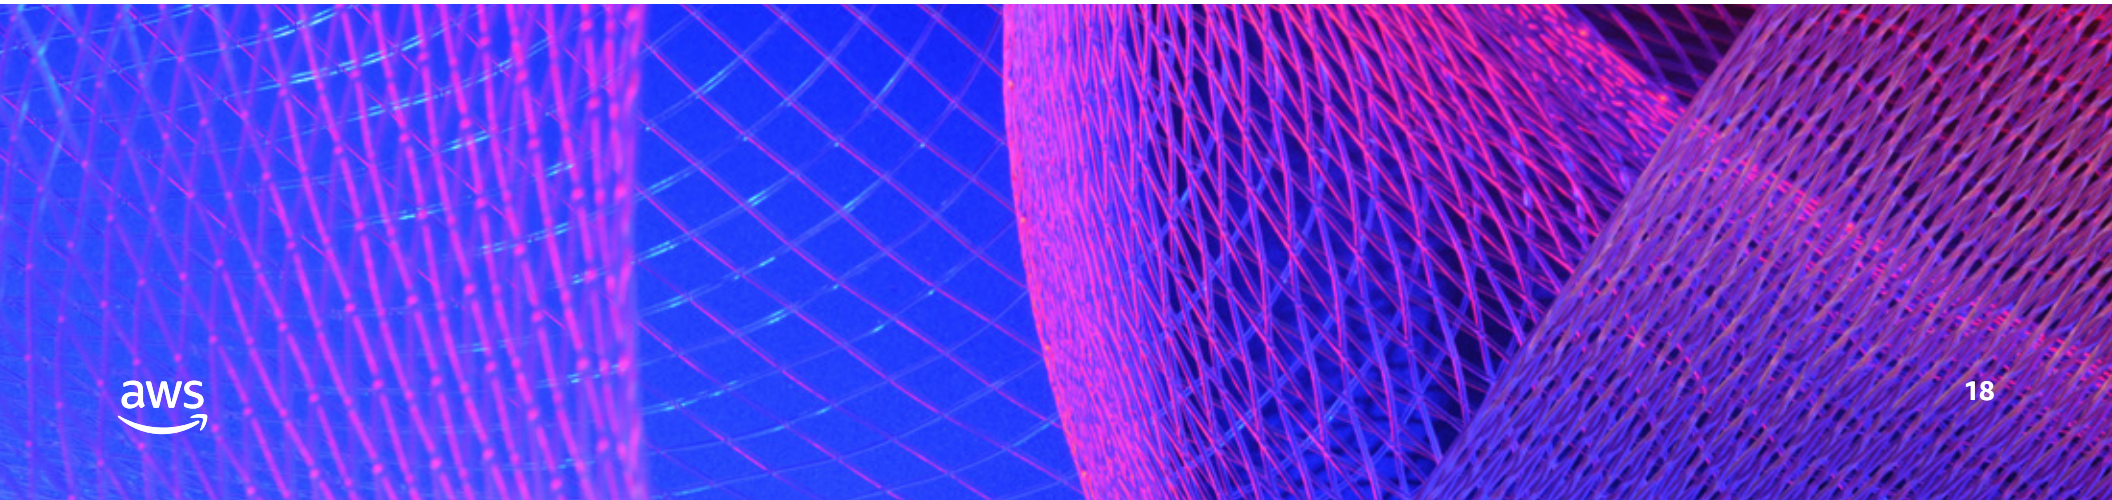

How AI21 Labs did it

AI21 Labs (AI21) is a leader in gen AI and LLMs. Initially, the company released two models: one with seven billion parameters and another with 178 billion parameters. But its team saw an opportunity to offer customers an LLM of 17 billion parameters to bridge the gap between the existing sizes.

AWS trained the FM in under 20 days using **Amazon SageMaker**—saving several weeks of time compared with AI21's previous training methods. "Because Amazon SageMaker handles node failures, restarts elegantly, and orchestrates large, distributed runs, the team working on pre-training the model can focus on core tasks," says Dan Padnos, vice president of platform at AI21.

"We have a really good relationship with the AWS team. Team members have gone deep into the technical details with us and collaborated on challenging tasks. Throughout the process, the AWS team has been creative and has had awareness about our challenges and goals."

Dan Padnos, VP of Platform, AI21 Labs

NEXT STEPS

Forging ahead with generative AI on AWS

No matter where your software company is in its AI adoption journey, you can take the next step with AWS solutions built on one of the most comprehensive cloud platforms and optimized for gen AI with high performance computing (HPC), security, and analytics. Empowered by a broad and deep set of AI services, over tens of thousands of organizations are running their AI workloads on AWS.

Gen AI can help you realize new business value within your organization. From reinventing customer experiences to enhancing productivity and accelerating growth, gen AI holds the power to help you transform not only your own business but your clients' operations, as well.

Artificial intelligence on AWS (by the numbers)

100,000+

customers using AWS for their AI workloads

20+YEARS

of building AI experiences at Amazon

$230M

committed to speeding creation of AI apps by startups

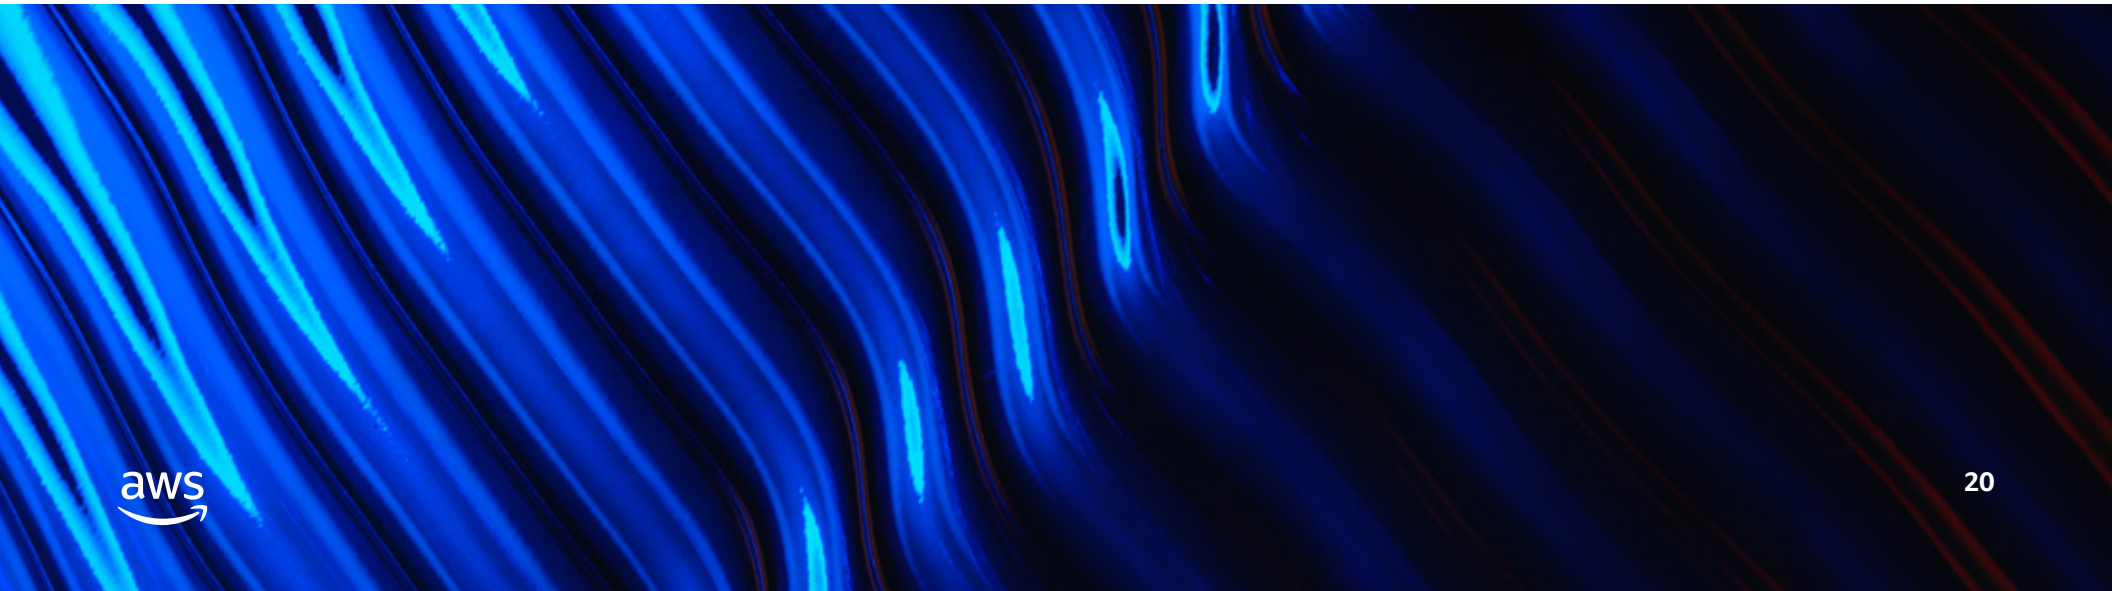

Enhance customer experiences with:

Chatbots and virtual assistants: Streamline customer self-service and reduce operational costs by automating responses for customer service queries through gen AI–powered chatbots, voice bots, and virtual assistants.

Agent assist and call analytics: Concisely summarize customer conversations to reduce the time agents and supervisors spend taking and reviewing notes or sharing context when transferring contacts. Analyze customer interactions to derive insights and monitor agent performance.

Hyper-personalization: Deliver better personalized experiences and increase customer engagement with personalized offerings and communications.

Boost employee productivity with:

Conversational search: Improve employee productivity by quickly and easily finding accurate information and summarizing content through a conversational interface.

Code generation: Accelerate application development with code suggestions based on the developer's comments and code.

Automated report generation: Automatically generate financial reports, summaries, and projections with gen AI, saving time and reducing errors.

AI-generated marketing content: Create engaging marketing content, such as blog posts, social media updates, or email newsletters, saving time and resources.

Sales: Generate personalized emails and messages based on prospect's profile and behavior, improving response rates. Generate sales scripts or talking points based on the customer's segment, industry, and product or service.

AI-generated sales content, guidance, and enablement: Generate personalized emails or messages based on a prospect's profile and behavior, improving response rates. Generate sales scripts or talking points based on the customer's segment, the industry, and the product or service.

Accelerate process optimization with:

IDP: Improve business operations by automatically extracting and summarizing data from documents and insights through gen AI–powered questioning and answering.

Data augmentation: Generate synthetic data to train ML models when the original dataset is small, imbalanced, or sensitive.

Supply chain optimization: Improve logistics and reduce costs by evaluating and optimizing different supply chain scenarios.

Solving the top artificial intelligence challenges

Many software companies have made investments and progress on their AI journeys and are exploring the possibilities of gen AI—only to find themselves hitting stopgaps along the way, worried that costs and complexities will grow too high as they progress. Throughout this eBook, we explored the steps to forge ahead and realize the full power of gen AI. To recap, let's look at the biggest challenges we identified along the way and brief recommendations on how your business can solve them.

To learn more about how gen AI can boost your software company's productivity, help build differentiated experiences, and encourage your teams to innovate faster, **visit the AWS generative AI page**.

Let's get started →

To learn more about how your software company can overcome obstacles and accelerate your AI journey, **visit the AWS AI resources hub**.

Challenge	Solution
Discouraging failures	Develop a fault-tolerant culture
Siloed, unprocessed data	Create a data strategy that includes data lakes
Finding the right business problems	Build blended teams that include both technical and domain experts
The AI skills gap	Adopt new organizational models, processes, and team management philosophies
Sustainably scaling beyond pilot projects	Leverage end-to-end tools like Amazon Bedrock and Amazon SageMaker to build and scale gen AI applications
Measuring the results	Forgo traditional ROI metrics in favor of the value tree model: agility, competitive advantage, and risk tolerance