



READY. SET. SCALE.

**A practical primer on scaling
AI for business value**

Accenture Applied Intelligence

About the authors



Fernando Lucini

Global Data Science and Machine Learning Engineering Lead, Accenture

Fernando Lucini is the Global Data Science and Machine Learning Engineering Lead for Accenture, managing the ongoing growth and effectiveness of their Data Scientists and ML engineers as well as their scaling methods and measurements of quality and effectiveness. He also leads Artificial Intelligence in the UK and Ireland. Fernando is a passionate and experienced senior leader with extensive experience in Artificial Intelligence and Machine Learning. Previously Fernando spent 18+ years in the enterprise software industry, creating technologies to automate and understand text, speech and video data and integrating these into business solutions for various Fortune 100 companies.



Sanjeev Vohra

Global Lead of Accenture Applied Intelligence

Sanjeev Vohra is the global lead of Accenture Applied Intelligence. A seasoned leader and respected innovator with over 18 years at Accenture, Sanjeev has served in a variety of leadership roles focused on large-scale digital transformation to create new growth opportunities. He is also a member of Accenture's Global Management Committee.

In his current role, Sanjeev helps clients across industries and markets evaluate, maximize and scale business value through the strategic application of data, advanced analytics, artificial intelligence and automation. He is responsible for the company's Applied Intelligence strategy, business growth, talent development, industry-specific solutions, strategic partnerships and V&A investments.



Table of Contents

- 01** Focus on value: Your business strategy is your AI strategy
- 02** Rethink work and get your people ready
- 03** Still having trust issues? Build responsibility into your AI
- 04** Productionize, and get ready to realize value
- 05** Get in touch




THINK BIG AND SCALE WITH INTENTION

3x

Strategic Scalers achieved nearly triple **the return from AI investments** compared to their non-scaling counterparts.

Why are so many businesses still struggling to realize the full value of their AI projects? The answer appears to lie in the gap between applying AI in pockets versus applying it at scale. Our latest Accenture Research, [AI: Built to Scale](#), indicates that Strategic Scalers—organizations that are scaling AI effectively—are seeing up to three times the return on their investments compared to non-scalers, or the ‘Proof of Concept (POC) group,’ who are conducting AI experiments and pilots but achieving a low scaling success rate and low return on their AI investments. Three out of four executives believe they will be out of business in five years if they can’t crack the code of scaling AI effectively. Companies need to get their heads around how to scale AI and fast—as global revenues from AI for enterprise applications are projected to grow by more than 50 percent between 2018 and 2025.¹






In practice, companies still find it difficult to make the transition from thinking about AI as a source of innovation to a critical source of business value.


There's a [state of paralysis](#) beyond the pilot.² Why? Until now, there hasn't been a proven blueprint for scaling, and organizations can fall into some common traps. First, companies don't have an AI roadmap or 'route to live'—the steps to take their AI project from POC to production, effectively and expediently. AI is different from 'traditional' software implementation projects, which companies are typically set up to deliver. Changing the status quo requires agility, openness to trying a new way of working and the ability to recognize when an idea works—and when it needs to be scrapped.

Second, the unfamiliar landscape of AI also means businesses can be tempted to fall back on their time-honored behaviors, reinventing the wheel and building from scratch. Big mistake. There are many proven, low-cost AI options to buy 'off the shelf' and start using right away. It is key to leverage what already exists, customize as needed for the organization



The difference in return on AI investments between companies in the Proof of Concept stage and Strategic Scalers equates to an average of US\$110 million.

and start proving the value of AI as the first step to successful scaling. But don't get bogged down in the technology. Be driven by the business strategy and vision, and let that dictate the AI approach. Focus on finding the right way of working that will allow AI to flourish, diversifying skills and talent beyond the data scientists. And get the [right governance approach](#) in place from the outset, with outcomes in mind.³



The following chapters will help you think through these critical success factors and understand how you can unlock a new wave of exponential value by scaling AI successfully.



01

FOCUS ON VALUE
YOUR BUSINESS
STRATEGY IS YOUR
AI STRATEGY

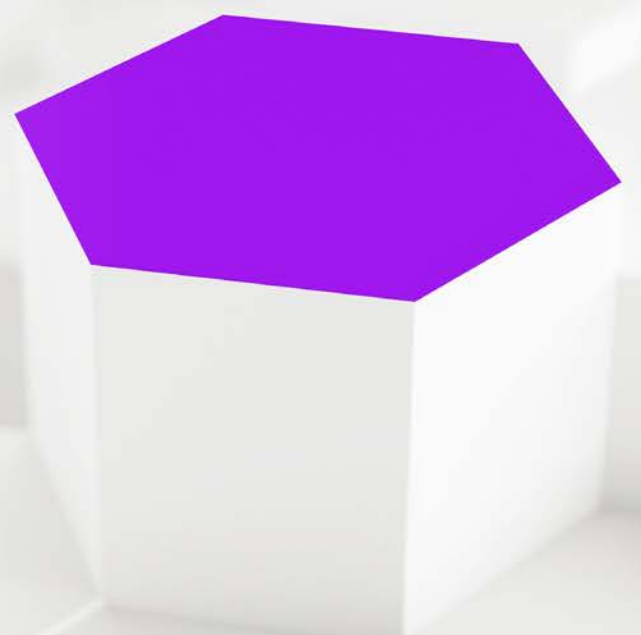


To scale AI successfully, get your ducks in a row early. That means:

- 01** Understanding what business value means to you
- 02** Translating that definition into a business strategy and
- 03** Focusing in on AI solutions that explicitly deliver on the most critical elements of that strategy.

Simple, right? If you have already defined value in your own unique context, you can harness AI to multiply that value—not just grow it marginally—charting a course that genuinely aligns with your business’s strategic priorities and delivers unprecedented returns. In our research, [AI: Built to Scale](#), Strategic Scalers understand this imperative, with more than 70 percent linking their AI ambitions explicitly to their overall business strategy.⁴

Decide what to focus on—and focus.



71%
of Strategic Scalers have
a **clearly defined strategy**
and operating model.



Look to the highest-level priorities.

It seems like more and more applications of AI are emerging each day. So, how do you determine which applications are going to deliver value, whatever that means for your specific context?

Finding true value starts with defining what really matters to a business and aligning the AI agenda to the highest-level strategic plans.

Ask yourself: What are the boardroom's short- and long-term priorities? How can AI help achieve the objectives of the C-suite such as organic growth, expansion into a new industry or development of new products?

For example, your company might look at 300-400 potential applications of AI ('use cases') as a starting point. They can't all be of equal value—or equally aligned to the overriding business strategy. A clear framework for identifying value should be used to qualify and prioritize AI use cases. To help get to that framework, you should think about these key questions:

- 1. What growth levers can you pull to deliver the most value?**
- 2. What is the potential business value in monetary terms?**
- 3. Will the use case be a differentiator in the marketplace, compared with those of competitors?**
- 4. Will it make our co-workers' lives simpler and better?**
- 5. Does the use case improve customer satisfaction and promote loyalty?**



Define value for today—with a vision for tomorrow.

While you need to look at the short-term return AI can create for your business, you also need to look at value—and therefore your high-level priorities—through a broader lens. Where is your organization headed in the ‘human-plus-machine’ era? What is the future of your industry? Will that change how you define value three to five years from now?

AI has the power to disrupt well beyond individual businesses. It is already blurring traditional industry boundaries, threatening legacy companies and giving agile new entrants the chance to make an impact, fast. Make sure you’re paying attention to what’s disrupting your industry already, how your world and the world at large are changing, and adjust your strategy, act boldly and invest to buy your way into the action. You may find yourself making different choices when you bring the macro into play.

For example, self-driving cars are dismantling the distinction between the car industry and [the insurance industry](#).⁵ New questions are emerging—

will the car industry require more safety features on self-driving cars as they evolve to keep people safe? Will the insurance industry be pushing for the same or different features to reduce claims? What about the effects of enhanced safety on vehicle design, insurance policies, emergency response, road network design, parking needs, car dealerships, petrol taxes and lawyers?

This is just one example of AI’s power to redefine industries. It shines a light on how your organization may need to change with the context of the larger-scale changes around you. We believe we can get under the surface of disruption—to understand how it is changing industries and how to capitalize on its opportunities. Our [Disruptability Index](#) can help you determine where your industry sits and how you can act now.⁶



93%

of executives say they know their industry will be disrupted at some point in the next five years, but **only 20% feel they're well prepared to address it.**⁷

Decide what value looks like for you—now and three years down the line. Don't be so focused on delivering for today that you aren't prepared for the next wave. Understand how AI is changing your industry and the world—and **have a plan to capitalize on it and stay relevant.**



Think practical, not provocative.

Leading AI organizations don't prioritize interesting or novel use cases that aren't aligned to the direction of their company, how their industry is developing or how society is evolving. Additionally, they tend to think practically about what will work in their operating conditions, ensuring that...

- The required data is readily available or easily sourced
- The use case can be integrated with existing tools and delivered to the front line
- The necessary skills and capabilities are in place and able to execute the plan
- There are data privacy or security risk protocols in place to mitigate unintended consequences
- There's an established method for calculating value and ROI



Take a portfolio view of your AI projects.

To be successful on your AI journey, **think about your AI projects as a portfolio of things you're trying to achieve.** This means thinking holistically about where you're headed and navigating the iterative nature of AI initiatives while remaining aligned to strategy and value. Scaling value relies on a formally defined AI roadmap which can help you deliver faster with more rigor and get to production more quickly.

The first step of the lifecycle is to create an 'idea pipeline'—and populate it with potential AI concepts that are yet to be tested for feasibility and value for your business. Shape, develop and investigate those ideas iteratively—but quickly—before a 'go/no go' decision. The ideas you generate may vary in terms of their potential to succeed, so having a holistic view of the collective success of your AI projects will be vital.

For example, data-driven technologies such as machine learning have a high risk/reward profile because development proceeds through a series of experiments. These experiments continually solve business problems by moving the engineering solution toward the optimal data pre-processing, learning algorithm and algorithm parameters for that specific use case.

Stay focused on your priorities while you learn.

At Accenture, we often find ourselves advising clients to spend at least 80 percent of their time designing AI that aligns very directly to their strategic priorities—and 20 percent or less on experimenting and learning for pure capacity building in the absence of a very clear target outcome. Experience tells us that when it comes to AI, even within that 80 percent, there's still a lot of experimentation!



Many of your experiments may not give you the outcome you think you needed. A solution is never guaranteed—and that's OK.

Therefore, assimilating AI into your business brings a new type of project execution risk, with only a portion of your ideas and experiments expected to go to production. But the good news is that following an AI roadmap, like the one here, helps qualify ideas quickly and effectively—so ideas that fail, fail fast and can be shelved with minimal investment before moving on to something else.

Strategic Scalers have mastered this approach. This group pilots more initiatives and successfully scales more often than their counterparts: They reported scaling 114 applications in the past three years, compared to just 53 for companies in the POC group.⁸





Underpin your AI strategy with a data strategy.

Every AI transformation journey starts with data. Nearly all AI Strategic Scalers (72%) agree that a [core data foundation](#) is an important success factor for scaling AI.⁹ More specifically, they understand the importance of having a data strategy—a design and intent that underpin what data is being captured, in what way, and for what purpose. The data strategy drives value as much as AI does.

And more data is not always better. In a world where data is proliferating and data begets more data, it can be tempting to gather more and more. Having a strong data strategy ensures you're curating the right data to deliver the desired outcome and then capturing its insights to fuel an AI strategy that delivers that outcome at speed and scale.

Once the data strategy is set, data can be mined to generate insights that help refine both the organization's strategy and the AI systems themselves. To really get the most out of this constant stream of data-driven insights,

you'll need to explicitly integrate 'feedback loops' into business decisions in an orchestrated way—for instance, to fine tune your business strategy and/or make necessary adjustments to your AI initiatives at the same time. This requires a new way of working: an agile, iterative approach to decision making—as well as AI development—with data at the core.

72%

of Strategic Scalers have a core data foundation, **compared to just 63% of their non-scaling counterparts** in the POC group.



Self-Assessment #1: Value and AI Strategy.

Now that you have read our empirical insights on defining value and an AI strategy, here are some practical questions to help you ‘self-calibrate’ and determine your readiness for AI at scale.

01 What are your organization’s strategic objectives and vision? How can AI help you achieve your desired outcomes?

a. What does ‘value’ mean to your company? What does it look like? And how do you currently calculate it?

b. Have you thought about what AI use case(s) will provide most value for you?

c. Are these use cases feasible? Do you have the necessary elements—data, technology, capabilities, etc.—within your company, or with partners, to deliver the use case(s)?

02 Do you have a data strategy? For example, are you able to find and use your data in ways that will help you achieve the objectives above? Is there any data you’re not using that you wish you could extract value from?

a. Are you able to reuse your valuable data and insights or do you recreate the wheel each time?

b. Are you mobilizing your data with internal resources or are you working with partners?

03 What part of your business (function), if any, is most AI-mature/ready? Are you already implementing AI in your organization? Are you running it within the business or in the technology function?


04 Do your current or planned AI use cases conform to your ethical framework? Are they aligned to a responsible purpose?

a. What are the ethical norms and regulatory considerations in your industry and/or market?

b. Or, are you operating in any markets that have launched or are preparing to launch their own AI policies, strategies and initiatives? If so, have you thought about what you can leverage for your organization?

02


**RETHINK WORK AND
GET YOUR
PEOPLE READY**





Challenge your assumptions about 'work' and get your people ready.

AI's disruptive nature means your old ways of working will need to change. Those who can successfully integrate AI into their culture and processes will be able to multiply value for businesses, employees and customers alike. Our [AI: Built To Scale](#) research confirms the correlation: Strategic Scalars are more likely than those in the Proof of Concept group to embed AI ownership and accountability into teams and ensure employees fully understand AI and how it relates to their roles.¹⁰





Word on the street.

As AI has moved into the adoption phase, urgent calls for reskilling and fears about job loss have taken center stage.

But we believe the majority of jobs will be reconfigured—not lost—as human-machine collaboration evolves. In fact, high-performing organizations will need big and diverse teams to scale AI successfully, with AI Strategic Scalers using twice the number of team members (about 200 people) for their AI projects compared to those in the POC stage.

Some roles will be totally new. For example, we have the ‘AI trainer’ who intervenes and manages AI technologies—helping systems ‘learn’ tasks and adjust conclusions—and helps us stay on the straight-and-narrow when it comes to ethics. As we learn new skills and delegate work to AI technologies, we can help the AI learn as well. It’s a team, just not as we know it.

Reimagine how work gets done.

The business of the future looks starkly different when your abilities can be multiplied with the use of machines. So how do you integrate AI into the way you work to maximize productivity? How do you reallocate workloads and reskill/upskill employees? It’s imperative businesses think about these questions to successfully integrate AI into their culture and processes.

Benefits across the business.

Companies from around the world and across industries are using AI to change the fabric of what they do and how they do it. Strategic Scalers are more likely to achieve a range of benefits including:

Leading organizations understand that the human + machine era will require them to break down and rebuild traditional notions of jobs and activities.



Start configuring the business of the future—now.


There are some practical steps you can take to start configuring business processes and the workforce to support AI at scale. Here are some key considerations to help you organize your business and improve agility in the way your people work.

01 Move from 'workforce' planning to 'work' planning

Break down traditional job roles, and look at which tasks and activities will be automated, which will require [human-machine collaboration](#), and how this might impact how people and teams intersect and interact.¹¹

02 Look seriously at new skilling

Get a clear view of the knowledge and skills you'll need to generate real value from human-machine collaboration. Look at your leadership, learning and recruitment programs, and invest in new ways to teach new things. For example, we put [60 percent of the money we save](#) from investments in AI into our training programs.¹²



03 Look at the big picture

What entirely new jobs—such as the ‘AI trainer’—can AI create in the organization? Are we prepared for those in the context of new markets, products and customer experiences? Operational roles may become insight-driven roles. For instance, when autonomous vehicles are introduced into society, the long-haul trucker may be reskilled to perform high-level technical work as an ‘in-cab systems manager.’ He or she might be monitoring systems and optimizing routes—potentially for multiple vehicles.

As we lean into human + machine collaboration, many human tasks will be augmented by AI. For example, AI can provide enhanced views of real-time data to help support decision making—without the decision making itself necessarily being offloaded to AI. It’s important to be clear about the right boundary, or process, for the organization when it comes to the split between the human and the machine—including how that boundary may shift as the organization’s AI maturity continues to change. Successful scaling relies on understanding how

the organizational chart will change with the upskilling and reskilling of people to be ‘data native’ and with new ways of delineating jobs and tasks.

Your workforce may be more ready than you think to adjust. [Our research](#) on the Future Workforce says so.¹³ Now it’s up to you to take action.

AI may be good for workers:

62%

of workers believe that **AI will have a positive impact on their jobs**, and 67% say it’s important to **develop skills** to work with intelligent machines.¹⁴



Establish the right talent mix.

It's no surprise that you need new kinds of talent to create AI products and services that deliver value. But beware of thinking data scientists are the only ones who matter when it comes to creating a route to 'go-live.' You also need data integration experts, business analysts, data engineers, and software engineers among others—and enough of them, in the right configuration.

In addition to the technical skills, it's important your team is interdisciplinary, bringing industry, business, design and governance expertise in the right ways and early on. These areas of knowledge might be easy to overlook, but they play a crucial role in creating successful AI applications. Think about it—we are creating AI to solve real-life problems, so having diverse specialists with expertise across numerous fields will help create a better AI solution.

Additionally, AI projects have a habit of generating [new and unexpected demand](#) for knowledge and skills—and having an interdisciplinary team can create a buffer, covering for any skills shortfalls you may encounter on your pathway to scaling AI.¹⁵

You need the right mix of talent to move from POC to production.

Strategic scalers build diverse teams to scale AI:

- Data integration experts
- Platform engineers/application specialists
- Data visualization experts
- Business analysts
- Data scientists
- Business unit leaders
- Data modelers

Look at your organizational set-up.

Along with establishing who gets the work done, it's important to revisit the 'how.' A priority should be to set up your organizational structure and governance strategy and establish the right connections within the business. Think about the kind of physical set-up that will help you achieve your business goals and integrate AI most effectively. For instance, do you need geographically dispersed business units and AI tools or a more centralized structure?

A [centralized organizational model may be the most effective](#), with Strategic Scalers saying they now use this approach.¹⁶ This model uses a centralized cross-functional AI group (sometimes called a Center of Excellence) to deliver projects to the various businesses, geographic units, and functions. This central group defines a standard way of working, governance model, and method of interaction for solution selection and implementation.

Another variant is 'hub-and-spoke,' a model that includes both a centralized cross functional AI group (i.e., the 'hub,' sometimes called a Center of Excellence) and separate autonomous AI teams (i.e., the 'spokes') that sit within business units. The hub provides a common set

of services and knowledge relevant to all business groups or functions, and sets common guidelines/operating principles, with some guidance from the business unit teams in the spokes. The spokes operate in a decentralized way to implement AI solutions relevant for their business units/functions.

Finally, a 'distributed' model also exists. Highly autonomous AI teams are housed within each business unit or function, with a delivery focus specific to that business unit or function. There are a few common standards, tools or operating principles imposed, with potentially varied practices across teams.

Highly distributed ways of working are the sign of mature technology, tools and usage patterns. Here, AI practitioners have considerable resources at their disposal and can be productive working in distributed but loosely connected groups with established ways of reusing and providing knowledge and expertise to the wider business.

Be guided by your business aspirations, and define a way of working that best supports those goals and your level of AI maturity.

Mind the gap.

There can be a gap between the CEO’s understanding of AI—what it can do and how—and what the people actually implementing the AI believe. The CEO’s perspective will naturally be influenced by the topline strategic intent of the company, what her peers are telling her, and what her long-term aspirations are for the organization. The AI leads doing the work might not always be aligned with the realities and focused goals of the C-suite—but they need to be! In fact, [our research](#) indicates that leadership’s limited understanding of AI’s potential can be one of the top challenges companies face when scaling AI. Strategic Scalers ‘mind the gap’—they reduce the distance between the goals and understanding of the C-suite and the practitioners when it comes to how AI can and should be applied to change the world, and their world.¹⁷

When you cut the distance between the C-suite and the AI practitioners, you improve your odds of delivering value.

Top challenges for scaling AI (Ranked top 5)

- ✓ Employee adoption of the solution **53%**
- ✓ Expected ROI from AI investments not adequate **52%**
- ✓ Inability to set up an organizational structure that supports continuous innovation... .. **48%**
- ✓ Foundational capabilities (data, technology) **47%**
- ✓ Failure to scale AI Pilots in the past **47%**
- ✓ Shortage of or difficult access to AI skills **46%**
- ✓ Lack of culture for change **43%**
- ✓ Poor understanding of AI potential within top management **42%**
- ✓ Inability to visualize the final outcome **42%**
- ✓ Lack of a clear roadmap **40%**
- ✓ Lack of budget **38%**




When it comes time to implement, look outside your organization for help.

We are now firmly in the 'era of implementation' with an explosion of investment in AI capabilities coming from well beyond Silicon Valley.¹⁸ These days, there are myriad tools which are proven, low-cost and academically rigorous. And there are varied and flexible ways to get your hands on AI: open source code, application programming interfaces (APIs), and small and medium-sized enterprise (SME) vendors to name just a few. As AI becomes mainstream, solution price points will also continue to drop.

AI start-ups are multiplying quickly as well. This means there's an ever-increasing pool of potential partners to help you on your AI journey. The figure below shows the number of AI start-ups in the market over time. Just look. In less than five years, the volume of machine learning AI start-ups has more than quadrupled.¹⁹

Now you can reuse, partner or buy to implement and scale AI capabilities before you even need to consider building new proprietary technologies in-house. Take advantage of what's out there for success at speed and scale.

Growth of AI Startups in the Market



Reuse before you buy. Buy before you build.

So how do you decide when to reuse, buy, partner or build? This is a full topic in and of itself (we [published an article](#) on it), but the simple answer is almost always reuse, buy or partner to take advantage of the investment other companies have already made—and get started quickly.²⁰

And ‘buying’ doesn’t need to mean purchasing an enterprise software product—rather, it’s a more fluid idea about paying for value and can include using a service, open source code or packages. We’ve helped clients, for example, move from buying speech-to-text software in its entirety to buying specific APIs that best solve the problem at hand—and pay for only what is used. By arbitraging use between the top vendors, we efficiently tackled the specific problem, sped up execution and avoided being stuck with the limitations of the software.

‘Buying’ can also mean purchasing an entire company with the capabilities to fulfill your future business needs. For example, in 2018, [Apple bought Silk Labs](#),²¹ a small start-up specialized in developing on-device machine learning (ML) services, and [it was reported](#) that the technology could potentially be used not only to improve Apple’s existing device solutions but also its ML tools for other applications.²² By acquiring companies that already have the people and capabilities in a specific area of interest, you can bypass steps and accelerate your scaling journey.

So you’re resolved to build—but are you doing it for the right reasons?

Sometimes, organizations decide to build AI in-house for the wrong reasons. For example, inexperienced teams may want to build from scratch as a way of learning on the job or trying to demonstrate their value—or they may believe they are the only ones who can get it ‘right’ for their companies. But don’t get caught in this trap. Build only because it makes business sense and the output will genuinely make you stand out.

Key questions to ask yourself to qualify the opportunity to build AI:

- Have we proven the value within an agile pilot? Prove the principle at small scale.
- Does someone in the business really need this, and is it a true differentiator? Be demand- and value- driven.
- Will they pay for it? Engage the business to generate investment commitment.

Self-Assessment #2: People and Capabilities.


Now that you have read our empirical insights on building your talent mix and capabilities, here are some practical questions to help you 'self-calibrate' and determine your readiness for AI at scale.

- 01** How much trapped value could be unlocked by moving POCs into production on a monthly basis?
- 02** Do you understand where the existing analytics tribes and data scientists are within your organization? Are they equipped to deliver against the business outcomes you're pursuing with AI or do you need to hire and create a more diverse talent pool?
- 03** Have you professionalized data science in your organization?

- 04** Do you have a route to 'live'? Do you have an established plan for productionizing to ensure that projects move past POC? Have you defined who is responsible for creating models and then operationalizing/productionizing the models? Are there gaps? Where are they? (Often the latter is a gap.)
- 05** Is your incentive structure set up to encourage the support of AI 'productionization'? Do you have a sense of the right breakdown between buy, partner, reuse, build? (Remember to build last, or only when it makes absolute business sense.)

03

**STILL HAVING
TRUST ISSUES?**





Still having trust issues? Build responsibility into your AI.

How do we learn to trust AI? Responsible AI builds trust and lays the foundation for successful scaling by taking a ‘human first’ approach—using technology to help people make better decisions, while keeping them firmly accountable through the right governance processes and technical steps. Our [AI:Built to Scale](#) research says responsibility is more than a ‘nice to have’—with AI Strategic Scalers significantly more likely to brief their employees clearly on how they tackle responsible AI.²³



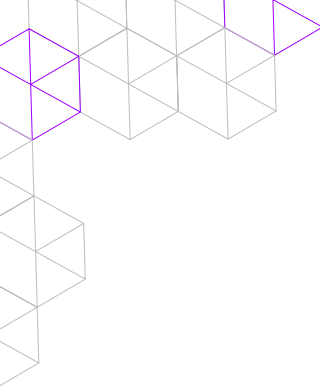
So you see the value in AI... but how do you trust it?

AI affords tremendous opportunities, from increasing efficiencies and improving outcomes, to reimagining industries altogether. Against this backdrop, it's easy to forget that AI's decisions also have a real bearing on people's lives, raising some big questions around ethics, trust, legality and responsibility. Enabling machines to make decisions may expose a business to significant risks, including reputational, employment/HR, data privacy, health and safety issues.

Enter: Responsible AI. It's a topic that's becoming pervasive in the media and a real consideration for clients in the public and private sectors.

What happens when a machine's decision turns out to be erroneous or unlawful? The potential fines and sanctions could threaten the business's commercial sustainability. And what about [other unintended consequences?](#)²⁴ AI has already shown it can be biased in ways that weren't anticipated and can hurt a brand's reputation. AmazonSM, for instance, had to scrap its AI-based recruiting tool that appeared to [show bias against women.](#)²⁵ And if need be, how does a human know when to intervene in a process driven by a machine?

In an Accenture [global executive survey](#), 60 percent of respondents believe that the adoption of AI is necessary to have a competitive advantage in a digital economy; however, 45 percent of our clients agree that not enough is understood about the unintended consequences of AI.²⁷



As AI increasingly permeates healthcare and government, the stakes get higher and potential unintended consequences could escalate. For instance, we will begin to trust AI tools to help doctors assess patients and determine if fast action is needed. [‘Explainable AI’](#) will be crucial to the longevity of these technologies.²⁶

We define Responsible AI as the practice of designing, building and deploying AI in a manner that empowers employees and businesses and fairly impacts customers and society—allowing companies to engender trust and scale AI with confidence.

What are unintended consequences?

Unintended consequences are the by-products of applying technology without accounting for the potential impact on a variety of unique communities and stakeholders—whether they are the intended audience for the technology or not.



Design trust into how you operate AI from the start. And from the top.

The Board of Directors needs to know what obligations it owes to its shareholders, employees and society at large, to ensure AI is deployed without unintended consequences.

The CEO might be asking, how can I be assured we have thought through AI's possible brand and PR risks? Meanwhile, the Chief Risk Officer and Chief Information Security Officer need to be thinking: If we deploy AI, how can we do it in a way that complies with data protection regulations? Creating a robust ethical underpinning for AI allows you to 'design out' legal and ethical concerns to the extent that it is possible.

However, it's not just about establishing the appropriate governance structures. It's also important to translate those ethical and legal frameworks into statistical concepts that can be unambiguously represented in software. In other words, there need to be clear and prescribed actions to support your ethical AI principles.

Putting the human at the center.

At Accenture, Responsible AI is more than just technology alone—it requires an interdisciplinary, innovation-friendly approach with humans at the core. The approach is rooted in our principles of TRUST:

- **Trustworthy**
Safe, honest and diverse
- **Reliable**
Enabling better decisions
- **Understandable**
Interpretable and transparent
- **Secure**
Strong data privacy and security
- **Teachable**
Human-centric in design



So, where to begin?

First, ensure considerations for AI are built into your core values and robust compliance processes. Then, you will need to implement specific technical guidelines to make sure that the AI systems are safe, transparent and accountable to protect your employees, clients, civilians, and other organizations.

Next, identify new and changing roles, and put the right training in place for technology specialists and your diverse team of experts to understand their new roles and remit. It's also critical to articulate a responsible business mission and conduct brand and public risk assessments to ensure your AI initiatives remain anchored to your core values, ethical guardrails, and accountability structure.

All of these elements are part of an innovation-friendly blueprint for Responsible AI that you can apply across functions and projects—allowing you to understand and manage the ethical implications of everything you do.



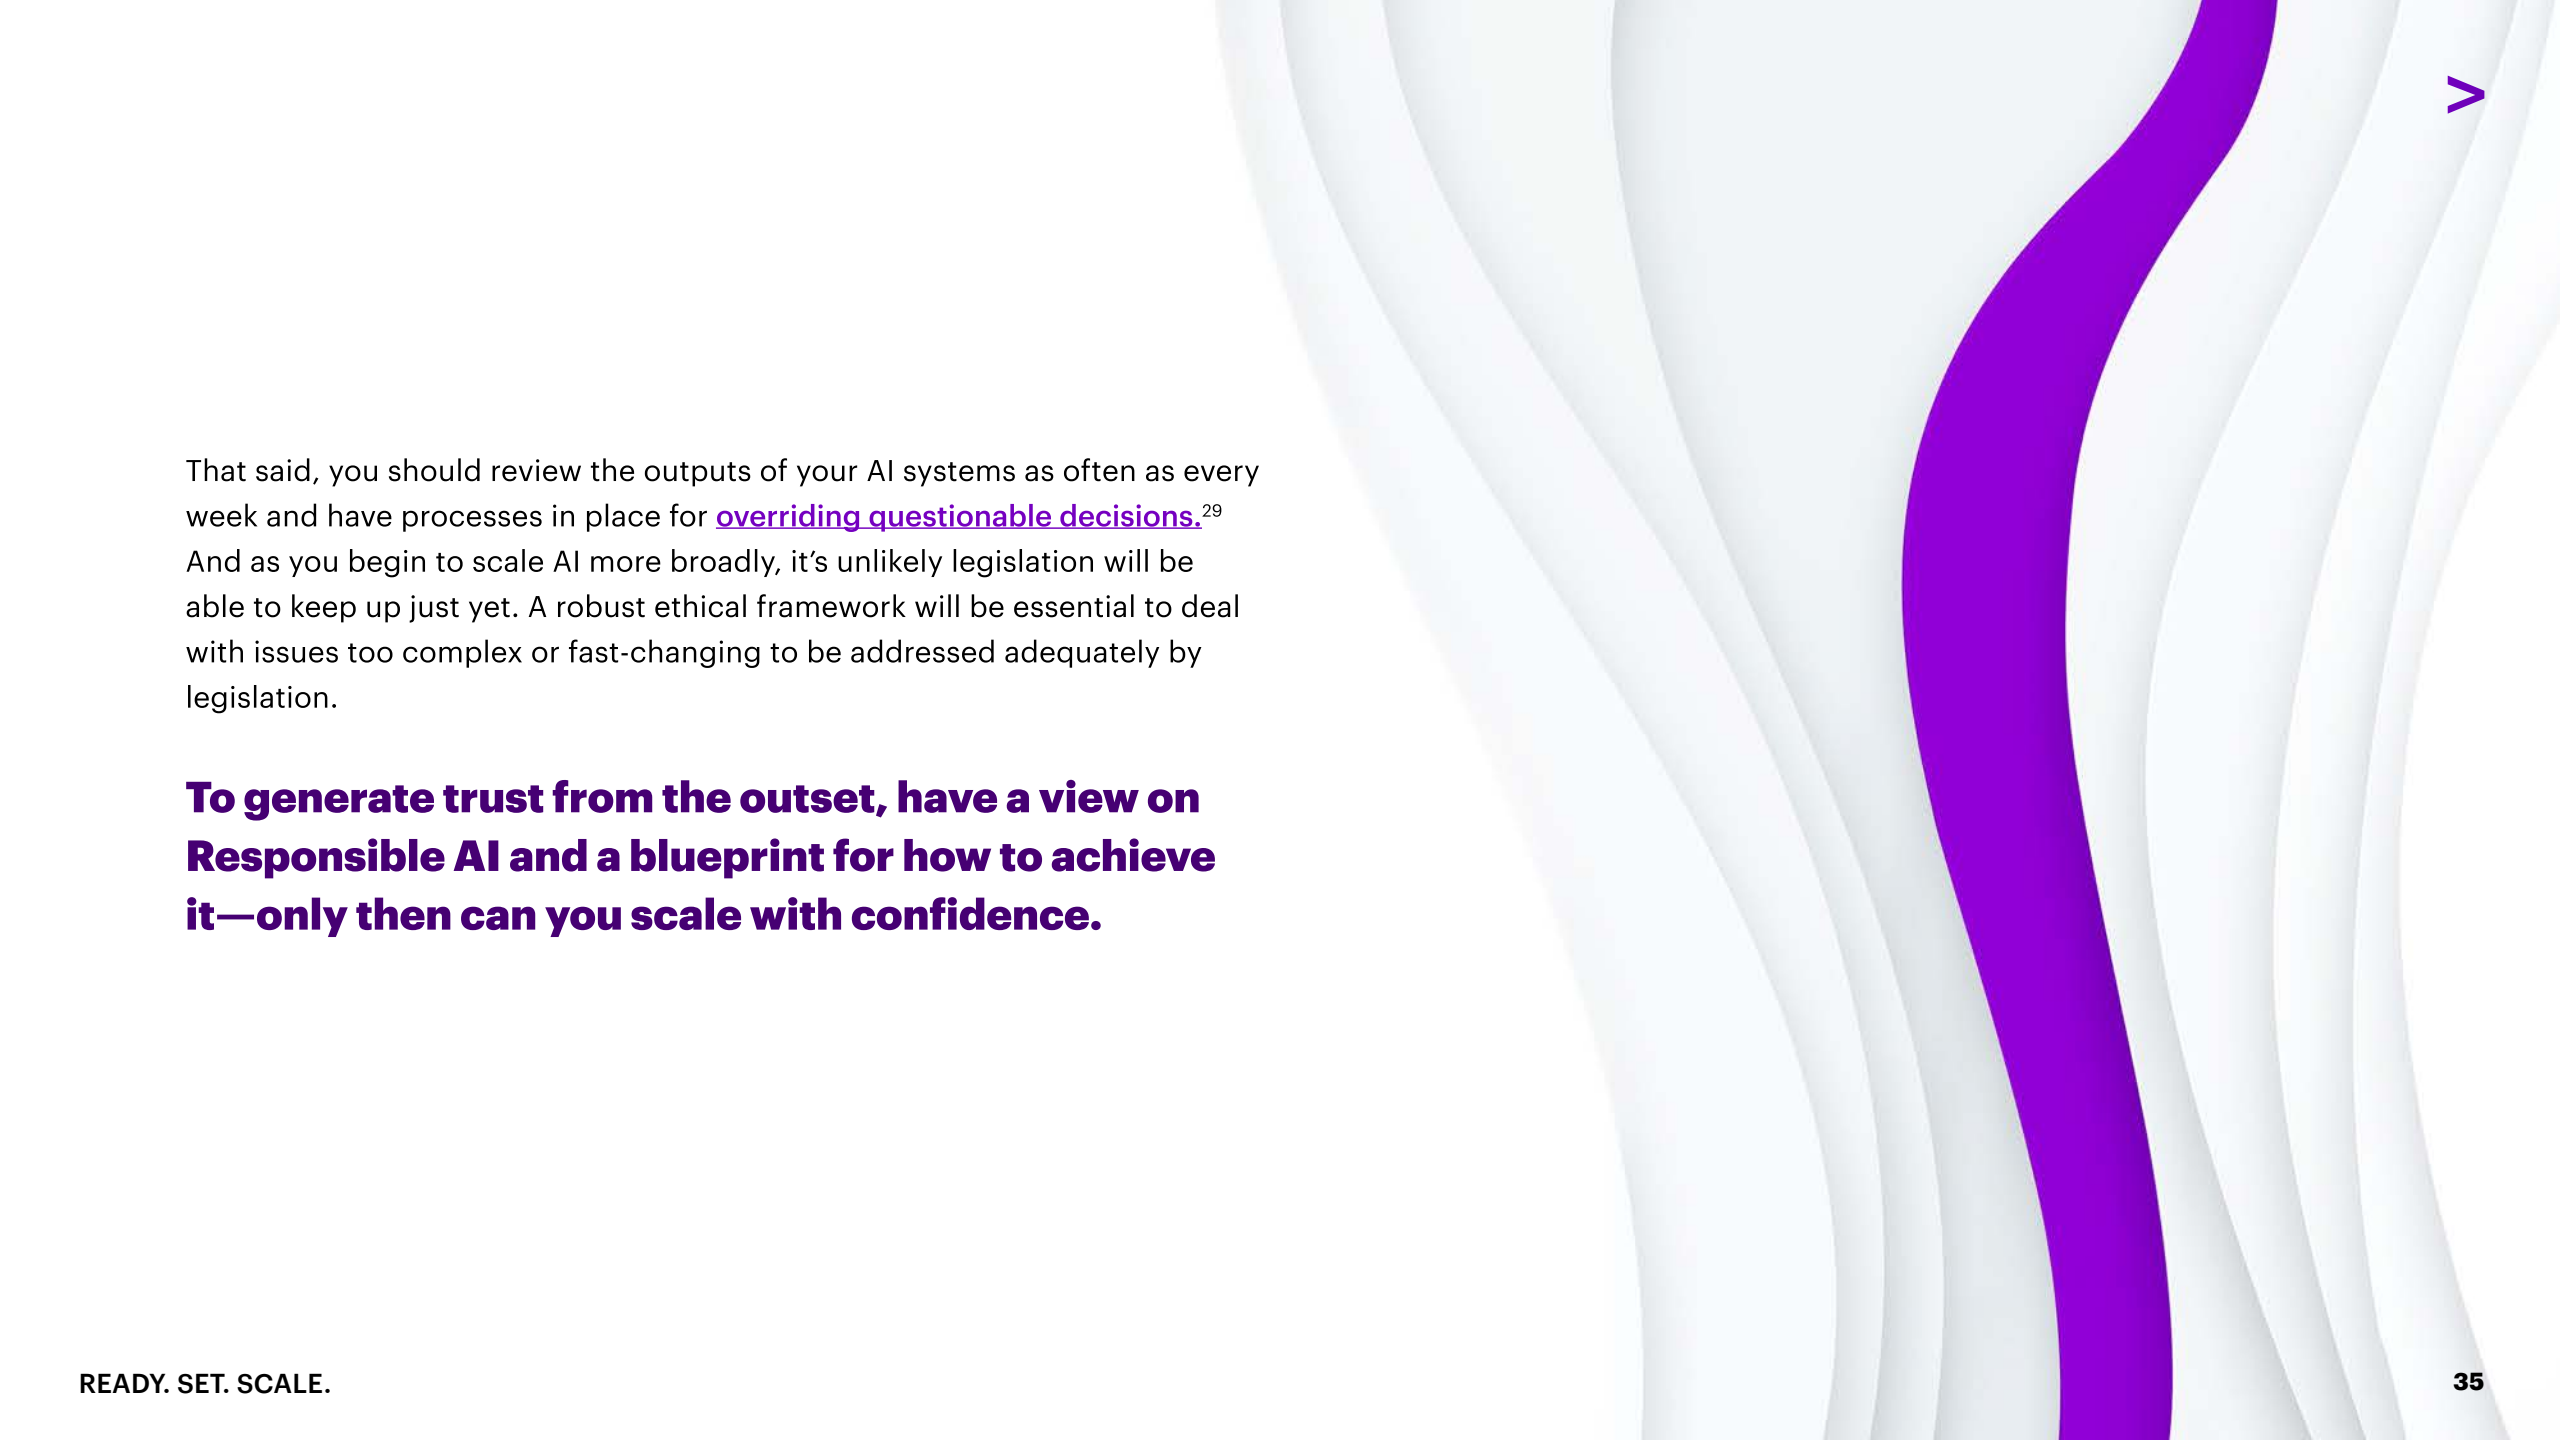
Police patrols versus fire wardens.

There are two main ways to **approach AI governance**: the police patrol and the fire warden.²⁸ Police patrols are the effect of AI governance rules applied from the top down. In this model, organizations monitor people and detect violations of the rules in order to enforce compliance. It seems like a straightforward option, but it tends to stifle innovation. Teams see AI governance as a barrier, rather than their shared responsibility.

The fire warden model is different. It embeds skills within teams to help them escalate issues that need attention—much like training fire wardens to raise the alarm and then carry out the necessary safety actions.

In general, we favor the latter approach for AI governance because it supports innovation and the agile development that is crucial to the competitiveness of today's businesses—and it can evolve more easily alongside fast-changing AI technology.

In an Accenture **global executive survey**, 88 percent of respondents indicated that they do not have confidence in AI-based decisions and outputs.³⁰ And 60 percent of respondents reported that they still require a human override of an AI system at least once a month.

The background features a series of light gray, wavy, vertical lines that create a sense of movement. A prominent, solid purple stripe runs vertically through the right side of the page, curving slightly. In the top right corner, there is a small, purple chevron symbol pointing to the right.

That said, you should review the outputs of your AI systems as often as every week and have processes in place for [overriding questionable decisions](#).²⁹ And as you begin to scale AI more broadly, it's unlikely legislation will be able to keep up just yet. A robust ethical framework will be essential to deal with issues too complex or fast-changing to be addressed adequately by legislation.

To generate trust from the outset, have a view on Responsible AI and a blueprint for how to achieve it—only then can you scale with confidence.



Put ethics at the core of AI development to build and retain trust.


Design in ethical frameworks when you're planning AI. We program algorithms to give us exactly what we have asked for, so we shouldn't be surprised when they do. And the problem is that simple algorithms treat all data as immutable, even data about our preferences, income and life situation. What can happen then, is that algorithms can trap people in their origins, history or a stereotype. These 'bad feedback loops' can lead to negative impacts on society.

The issues mentioned are not inherent to machine learning algorithms themselves. Instead, issues arise from the way they interact with society and the unintended consequences that can result from those interactions. As such, putting the ethical implications at the heart of the development of each new algorithm is vital.

One way to ensure this is by embracing public health models of governance, which treat issues as indicative of underlying drivers, rather

than problems to be solved per se. Another would be to ensure algorithms can be adapted more readily to newer or better data, in ways that do not exaggerate historical patterns. We see this every day in the way that AI at Spotify™ or Amazon quickly adapts recommendations to our latest searches.

Finally, targeted research identifying individual problems and solutions is critical to the success of any effort to create more ethical AI. We need to see more resources—and more senior leadership attention—directed at making sure algorithms do not have negative impacts on individuals or society. Just as data privacy and cyber security have moved from departmental to board-level issues, responsible governance of AI must be quickly elevated in importance by all organizations that use it.



Trade your black box for a glass box.

With a recent upsurge in machine learning, especially deep learning, it can feel as if AI is working in a ‘black box.’ But when you’re on the receiving end of an AI-generated decision, how do you know how it was made or how to challenge it? For example, if AI is deciding whether or not to grant you a loan, how did it make that decision, and how would you dispute any of its conclusions?

In light of such issues, perhaps the greatest barrier to AI adoption among consumers will be building—and sustaining—trust in AI. This is where transparency and [explainable AI](#) come into play.³¹

Be open with consumers about when they are interacting with an AI capability. Be able to show how decisions are made. Create strong frameworks for human decision override and review—especially where it most impacts the consumer’s experience. And edit AI’s outcomes frequently

We will only begin to understand and trust AI when ‘explainability’ takes it out of the ‘black box’ and demystifies its decisions.

To meet the requirements around compliance, accountability and transparency coupled with the adoption of AI models, financial institutions will be required to explain the workings of their AI models. For example, [the Explainable AI \(XAI\) research project](#) run by the Defense Advanced Research Project Agency (DARPA) is developing a suite of ML techniques that can produce more explainable models—enabling humans to interpret, trust and manage AI systems.³²

Self-Assessment #3: AI Governance and Responsible AI.


Now that you have read our empirical insights on AI governance and Responsible AI, here are some practical questions to help you 'self-calibrate' and determine your readiness for AI at scale.

- 01** What's your company definition of fairness and bias? How do you translate that definition and your code of ethics into implementation standards (turning talk to action)?
- 02** Do you have an ethical data usage manifesto? Do your customers know if/that you're ethically responsible with their data and yours? Do you have regulation and compliance capabilities in place to manage data and AI technology?
- 03** What is your process for mitigating biases in the data you use to power AI initiatives?

- 04** Have you begun or are you planning to begin embedding mechanisms for achieving fairness, transparency, and accountability into your design, development, deployment and monitoring of AI systems? Who is accountable for the decisions made by AI systems?
- 05** Have you updated your risk frameworks to incorporate contingency plans for incorrect outcomes associated with a range of inputs, including the use of inaccurate historic data and lineage? Is there a formal process in place to mitigate unintended consequences of AI decisions?

04

**PRODUCTIONIZE
AND GET READY
TO REALIZE VALUE**





Scaling value is about understanding how to move from pilot to production; getting your data strategy in place to drive real-time strategic actions; and establishing the right talent mix, operating model and governance framework. Those who succeed will reap the rewards. And those who fail may find their businesses fall by the wayside (75 percent of executives believe they will be out of business in five years if they cannot scale AI effectively).³³

AI is no longer a 'nice to have' or a set of cool tools to impress management. AI and data strategies are becoming the very core of business, and all the while it's becoming easier and cheaper to get your hands on the technology. The time to act is now.



There's a lot to think about—and a strong business case to get started quickly. In this primer, we have asked some questions and provided some insights on what it takes to scale AI effectively and move beyond proofs of concept to production. But how does it all come together in practice—and what concrete steps can you take to realize value quickly?

The final chapter of this primer is our AI Roadmap, a start-to-end model we use with our clients to help them realize and multiply value from their AI projects. It details an AI use case's route to live, which includes defining value and formulating a solid AI strategy; bringing together the right AI capabilities; thinking about the optimal talent mix; and getting the appropriate governance and ethical parameters in place. But it doesn't just end there. It lays the path for how to multiply value from the use case through continuous engineering, optimization and the extension of the feature to new use cases.

We invite you to go through our roadmap and evaluate how you're approaching your AI projects. Stop at each checkpoint and ask yourself the flagged questions to make sure you're setting yourself up for success—with your data, your people, your infrastructure and your organization at large. Whether you've been in proofs of concept or are already starting to scale AI, be assured that there are concrete steps you can take to realize even more value from your AI initiatives.

So are you ready? Because now's the time to get set—and scale.

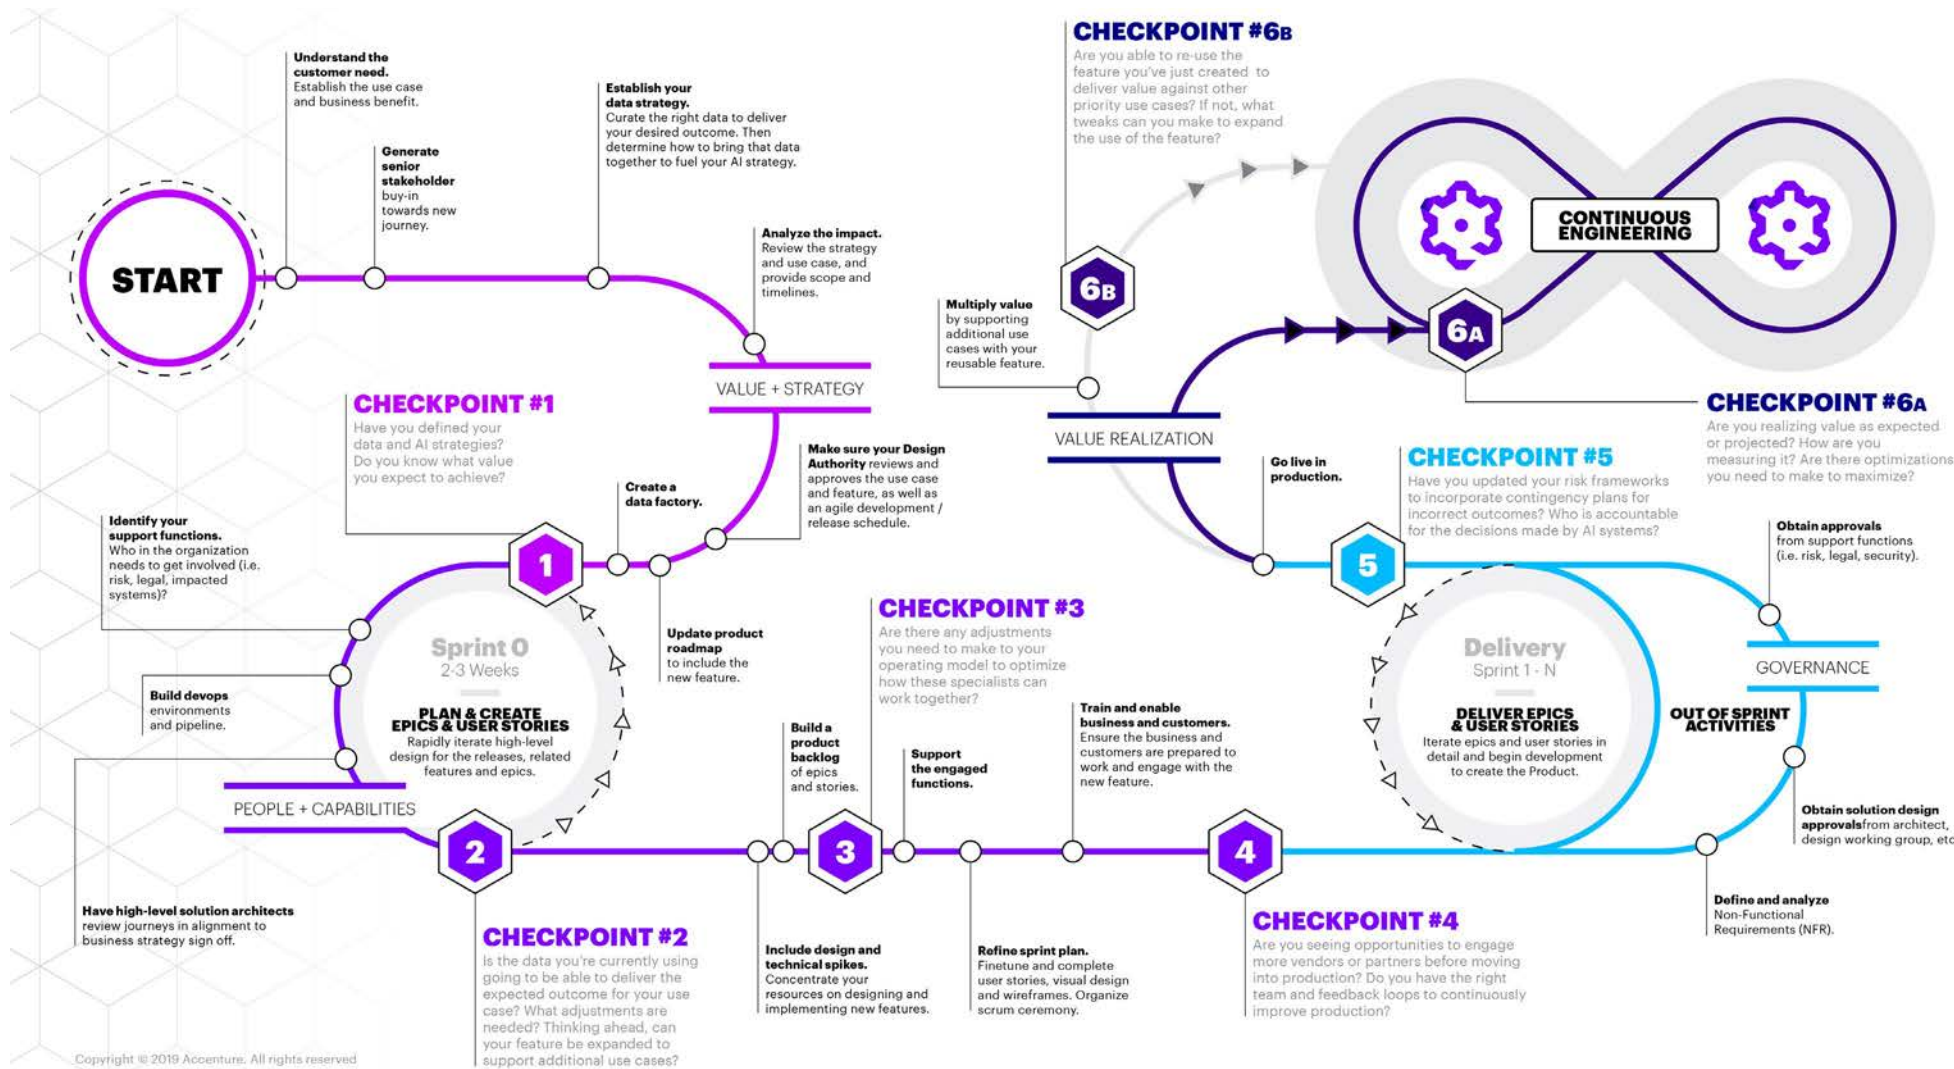
Recap: Are you built for AI at scale?

Have you defined what 'value' means to your business and which AI use cases to prioritize to deliver on that defined value? Do you have a data strategy in place, knowing data is a critical enabler of your AI technologies?

1. Do your AI teams understand the business goals you are trying to achieve with AI—and are they set up to move to production (and avoid being trapped in POCs)?
2. Do you have ethical frameworks and a way to course correct should the AI make erroneous, biased or unlawful decisions?



AI roadmap - The journey to live.




Copyright © 2019 Accenture. All rights reserved

To scale effectively, run an unbreakable thread that traces the critical path to production through all of these highly connected elements. Only then can you amplify value.

05

GET IN TOUCH





Contact the authors to find out more about how to scale AI no matter where you are on your AI journey—and see how value can be multiplied for your business.

Join the conversation.

[@AccentureAI](#)

www.linkedin.com/company/accentureai/

www.accenture.com/appliedintelligence

Authors.

Fernando Lucini

Global Data Science and Machine Learning
Engineering Lead, Accenture

Sanjeev Vohra

Global Lead of Accenture Applied Intelligence



References.

¹Accenture, “AI: Built to Scale,” November, 14, 2019. <https://www.accenture.com/us-en/insights/artificial-intelligence/ai-investments>

²Accenture, “Overcoming data paralysis when preparing for AI,” February 4, 2019. <https://www.accenture.com/us-en/insights/digital/preparing-data-ai>

³Accenture, “An AI governance approach to support innovation,” February 8, 2019. <https://www.accenture.com/us-en/insights/artificial-intelligence/ai-governance-support-innovation>

⁴Accenture, “AI: Built to Scale,” November, 14, 2019. <https://www.accenture.com/us-en/insights/artificial-intelligence/ai-investments>

⁵Accenture, “Insuring autonomous vehicles,” <https://www.accenture.com/us-en/insight-autonomous-vehicles-opportunity-insurers>

⁶Accenture, “Disruption need not be an enigma,” February 26, 2018. <https://www.accenture.com/us-en/insight-leading-new-disruptability-index>

⁷Accenture, “Disruption need not be an enigma,” February 26, 2018. <https://www.accenture.com/us-en/insight-leading-new-disruptability-index>

⁸Accenture, “AI: Built to Scale,” November, 14, 2019. <https://www.accenture.com/us-en/insights/artificial-intelligence/ai-investments>

⁹Accenture, “AI: Built to Scale,” November, 14, 2019. <https://www.accenture.com/us-en/insights/artificial-intelligence/ai-investments>

¹⁰Accenture, “AI: Built to Scale,” November, 14, 2019. <https://www.accenture.com/>

[us-en/insights/artificial-intelligence/ai-investments](https://www.accenture.com/us-en/insights/artificial-intelligence/ai-investments)

¹¹Accenture, “Empower the human+ worker,” February 7, 2019. <https://www.accenture.com/us-en/insights/technology/future-of-work>

¹²Accenture, “Reworking the revolution,” <https://www.accenture.com/gb-en/company-reworking-the-revolution-future-workforce>

¹³Accenture, “Reworking the revolution,” <https://www.accenture.com/gb-en/company-reworking-the-revolution-future-workforce>

¹⁴Accenture, “The big disconnect: AI, leaders and the workforce,” July 12, 2018. <https://www.accenture.com/us-en/insights/future-workforce/big-disconnect-ai-leaders-workforce>

¹⁵Accenture, “New, unexpected skills you need for AI success,” August 22, 2019. <https://www.accenture.com/us-en/insights/artificial-intelligence/new-unexpected-skills-ai-success>

¹⁶Accenture, “AI: Built to Scale,” November, 14, 2019. <https://www.accenture.com/us-en/insights/artificial-intelligence/ai-investments>

¹⁷Accenture, “AI: Built to Scale,” November, 14, 2019. <https://www.accenture.com/us-en/insights/artificial-intelligence/ai-investments>

¹⁸Dr. Kai-fu Lee, AI Superpowers: China, Silicon Valley and the New World Order. Houghton Mifflin Harcourt, 2018.

¹⁹Venture Scanner, “Artificial intelligence startup highlights Q2 2018.” <https://>



References.

www.venturescanner.com/2018/08/30/artificial-intelligence-startup-highlights-q2-2018/

²⁰Accenture, “What’s the best AI strategy: Build, buy, partner?,” May 6, 2019. <https://www.accenture.com/us-en/insights/digital/best-ai-strategy-build-buy-partner>

²¹The Verge, “Apple reportedly buys AI startup with privacy-conscious approach,” November 21, 2018. <https://www.theverge.com/2018/11/21/18106192/apple-privacy-ai-silk-labs-acquisition>

²²Venture Beat, “Apple reportedly bought Silk Labs, a maker of privacy-focused AI for devices,” November 21, 2018. <https://venturebeat.com/2018/11/21/apple-reportedly-bought-silk-labs-a-maker-of-privacy-focused-ai-for-devices/>

²³Accenture, “AI: Built to Scale,” November, 14, 2019. <https://www.accenture.com/us-en/insights/artificial-intelligence/ai-investments>

²⁴Accenture, “How to stop AI from reinforcing biases,” August 2, 2019. <https://www.accenture.com/us-en/insights/artificial-intelligence/stop-ai-reinforcing-biases>

²⁵Reuters, “Amazon scraps secret AI recruiting tool that showed bias against women,” October 9, 2018. <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G>

²⁶Accenture, “Explainable AI,” September 21, 2018. <https://www.accenture.com/us-en/insights/artificial-intelligence/explainable-ai>

²⁷Accenture, “Critical mass: Managing AI’s unstoppable progress,” September 17, 2018. <https://www.accenture.com/us-en/insights/digital/critical-mass-managing-ai-unstoppable-progress>

²⁸Accenture, “An AI governance approach to support innovation,” February 8, 2019. <https://www.accenture.com/us-en/insights/artificial-intelligence/ai-governance-support-innovation>

²⁹Accenture, “Critical mass: Managing AI’s unstoppable progress,” September 17, 2018. <https://www.accenture.com/us-en/insights/digital/critical-mass-managing-ai-unstoppable-progress>

³⁰Accenture, “Critical mass: Managing AI’s unstoppable progress,” September 17, 2018. <https://www.accenture.com/us-en/insights/digital/critical-mass-managing-ai-unstoppable-progress>

³¹Accenture, “Explainable AI,” September 21, 2018. <https://www.accenture.com/us-en/insights/artificial-intelligence/explainable-ai>

³²DARPA, “Explainable Artificial Intelligence (XAI).” <https://www.darpa.mil/program/explainable-artificial-intelligence>


³³Accenture, “AI: Built to Scale,” November, 14, 2019. <https://www.accenture.com/us-en/insights/artificial-intelligence/ai-investments>



About Applied Intelligence.

Applied Intelligence is Accenture's approach to scaling AI for our clients. We embed AI-powered data, analytics and automation capabilities into business workflows to accelerate time to value. Our expertise in defining end-to-end strategy, combined with deep data infrastructure capabilities, cognitive services and industrialized accelerators help smooth clients' path to AI adoption, extending human capabilities and supporting clients in scaling AI responsibly. Recognized as a leader by industry analysts, we collaborate with a powerful global alliance, innovation and delivery network to help clients deploy and scale AI within any market and industry.

Follow [@AccentureAI](#) and visit [accenture.com/appliedintelligence](https://www.accenture.com/appliedintelligence).




accenture

