

How to think—and act—like an AI-native business

Existing businesses can't become AI natives themselves. But they do need to stay vigilant as these AI upstarts seize new market opportunities. By actively studying how AI-native businesses put AI into the core of their operations and technology, established companies can reap the benefits of thinking and acting like their newest competitors.



Remember Monster.com? Back in the late 1990s, the emergence of the online job site was a godsend for job seekers everywhere who happily gave up their newspaper classified ad subscriptions for the new, ink-free web platform.

Within a few years, however, along came a new set of disruptors—from Indeed to LinkedIn to Glassdoor—each of which offered progressively more innovative features.

Soon, however, we may be asking a different question: Remember online job boards?

With the emergence of generative AI, we may soon experience a whole new approach to finding a job—or getting a date, mapping a route or booking a hotel. It seems like just yesterday these activities morphed into a new normal. Now, once again, they're up for disruption.

Imagine, for instance, an AI-powered career coach. Rather than searching for a job, candidates would tell it their skills and goals, even their ideal company culture. In return, beyond offering up matches, it would analyze hidden patterns in job postings, news articles and social data to uncover opportunities at companies that hadn't even advertised yet. Better yet, it would generate tailored cover letters, conduct practice interviews, even suggest negotiation tactics based on the potential employer. Who would make this happen? Most likely an AI-native business.

While AI-native businesses as we've defined them don't exist yet, they soon will. These AI-first upstarts will insert AI (and increasingly generative AI) into the fabric of everything they do. And because their whole mindset is centered on AI-driven capabilities, the technology will drive how they think and do things.

Unburdened by the constraints of legacy systems and entrenched approaches, AI-native businesses will see AI not as a tool to be bolted on, but as the fundamental building block of their operations.

This pure AI mindset will allow them to capitalize on generative AI's strengths, from the tailored experiences customers crave, to blazingly fast internal processes, to business models that, to a traditional business's eye, appear inside-out and upside-down.

Like the digital-native businesses before them (Uber, Netflix, Venmo), AI-native businesses will also change how consumers behave. Consumers will begin to choose products differently, seek information differently and assess options differently, ultimately [changing the way in which they interact with the digital world](#).


And that's why traditional businesses need to pay attention. As consumers change their behaviors, businesses will need to accommodate these new ways of selecting, engaging and transacting. These changes will also open up opportunities for AI-native businesses to infiltrate enterprise functions, both collaboratively and competitively:

- Existing companies might partner with AI-native businesses for specialized solutions, filling crucial gaps with capabilities they can't build in-house.
- They might also compete with AI-native businesses that offer products and services that completely upend established processes. Consider the AI-powered writing assistant that challenges traditional advertising agencies, or the predictive logistics platform that outperforms legacy supply chain management tools.


To execute newly hatched business models, organizations will need new governance policies, organizational structures and operating processes to deal with massively accelerated, dynamic and autonomous operations. And—vitality—they'll also need to rethink their technology infrastructure to power the new AI-driven capabilities they offer.

This is where thinking like an AI-native business comes in. [Existing companies can't become AI-native businesses—that's logically impossible](#). But they can and should learn from how AI-native businesses operate. By actively studying these newcomers' innovative approaches, understanding their technology stack and adapting their own AI strategies, established companies can start to think—and even act—like the disruptive businesses that will put AI at the heart of all they do.

Jump to a section:



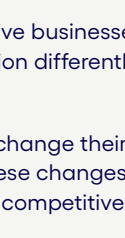
Widescale disruption in two years: our research reveals when real change will start



From niche specialists to game-changers: three profiles of AI-native businesses



Tech lessons from AI-native businesses: get ready for the new AI tech stack



Build, partner, utilize: strategic considerations for businesses in an AI-native world



Widescale disruption in two years

Our research reveals when real change will start

It's not today but soon that these levels of generative AI disruption will gain a foothold. To better understand how generative AI will change the future of work, we partnered with Oxford Economics to create an economic model analyzing the technology's impact on business productivity and the workforce itself. (For the full study, see our recent report [New work, new world](#).)

Based upon our findings on how quickly businesses will adopt generative AI, and our analysis of past technology advances, we've determined that generative AI adoption will follow an S-curve trajectory: a gradual rise, to a dramatic spike, to a plateau in which the technology becomes refined and pervasive. That dramatic spike will take place over the next decade, rolling out over three key phases (see Figure 1).

Fast uptake in the next decade

Adoption of generative AI could move quickly in the decade ahead.

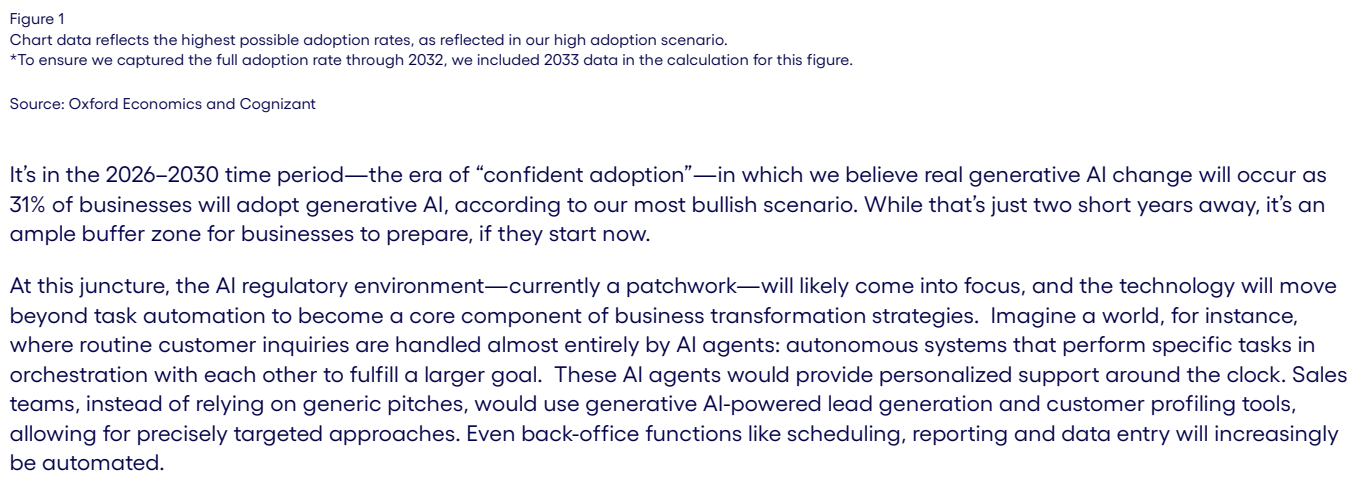


Figure 1
Chart data reflects the highest possible adoption rates, as reflected in our high adoption scenario.
*To ensure we captured the full adoption rate through 2032, we included 2033 data in the calculation for this figure.

Source: Oxford Economics and Cognizant

It's in the 2026–2030 time period—the era of “confident adoption”—in which we believe real generative AI change will occur as 31% of businesses will adopt generative AI, according to our most bullish scenario. While that's just two short years away, it's an ample buffer zone for businesses to prepare, if they start now.


At this juncture, the AI regulatory environment—currently a patchwork—will likely come into focus, and the technology will move beyond task automation to become a core component of business transformation strategies. Imagine a world, for instance, where routine customer inquiries are handled almost entirely by AI agents: autonomous systems that perform specific tasks in orchestration with each other to fulfill a larger goal. These AI agents would provide personalized support around the clock. Sales teams, instead of relying on generic pitch decks, would use generative AI-powered lead generation and customer profiling tools, allowing for precisely targeted approaches. Even back-office functions like scheduling, reporting and data entry will increasingly be automated.

But beyond imagining this world, our study brings many of these changes to light. As part of our analysis, we assigned “exposure scores” to 1,000 jobs currently being done by the US workforce.

These scores don't reflect the percent of workers who will be out of a job or their chance of losing a job. Rather, it's the maximum percent of job tasks that could theoretically be automated or assisted by generative AI by the year 2032, weighted by the relative importance of those tasks.

A look at our exposure scores reveals occupation groups that are due for the most change with the emergence of generative AI—and represent prime areas for AI natives to both enter new markets and emerge as lean businesses. When looking at the areas above, for instance, customer service representatives will see exposure scores of 63.7%, sales representatives 65.6% and office and administrative workers 85.9%, all by 2032.

This is why—starting in the next two years when confident adoption of generative AI begins—forward-thinking companies will need to have a robust technological infrastructure, a strategy in mind for redefining the business model and a culture that encourages continuous innovation. For that, we can turn to AI-native businesses, which can offer lessons on how and where to make these important shifts.



From niche specialists to game-changers

Three profiles of AI-native businesses

Niche specialists

Precise problem solvers



Platform builders

Foundation architects



Game changers

Industry disruptors



In our view, AI-native businesses will come in three main forms, each with the potential to reshape established industries. Niche specialists—similar to the career coach described earlier in this report—will focus on solving specific pain points with laser-like precision. An AI tool might work beside a skilled doctor, for instance, assisting in diagnosing rare diseases with an accuracy that surpasses the efforts of entire medical teams. Or, a manufacturing tool might be so tuned to the rhythm of an assembly line that it can predict impending breakdowns, preventing costly downtime and maximizing efficiency.

These purpose-built agents will likely be the first area where we'll see AI natives emerge. But that's just on the front end: For all three types of AI natives, much of what goes on behind the scenes will also be done by AI agents that collaborate on orchestrating tasks to get work done. [These armies of internal AI agents will become integral team members, reducing the need for large workforces and allowing for rapid decision-making](#). (For more on this topic, see Cognizant CTO of AI Babak Hodjat's [recently published article](#), “AI and agents.”)

Platform builders, meanwhile, will function like architects, crafting the foundational technologies that enable other businesses to harness the power of AI. These innovators will build ecosystems designed to elevate the employee and customer experience. Some platforms will focus on anticipating customer needs, offering personalized recommendations and solutions before they're even asked. Others will focus on decision-making, transforming raw numbers into actionable insights, moving far beyond simple dashboards.

But perhaps the most impactful of all are the game changers. These visionaries won't see boundaries—they'll only see potential. They will wield AI to disrupt entire industries by defying traditional limitations: providing affordable healthcare to remote communities, reinventing financial investments with AI-driven algorithms, or using the power of generative AI to unlock the full potential of personalized education accessible to everyone.

Let's explore three hypothetical examples of how each type of AI-native business could reshape everyday experiences and disrupt how work gets done:

Niche specialist

A prescriptive supply chain oracle

This AI-native business would flip traditional supply chain planning on its head. Instead of addressing problems reactively, it would anticipate slowdowns and risks and suggest ways to avoid them.

The platform would ingest a vast and seemingly chaotic stream of data: weather patterns, port congestion, raw material prices, local events, political instability and even social media sentiment. It would then uncover correlations invisible to human analysts, like a spike in online discussions about a specific ingredient, signaling an impending shortage in time for businesses to source alternatives.

It would also proactively suggest opportunities, identifying surplus inventory or transport routes newly opened due to shifting geopolitical situations. Different stakeholders would get tailored reports: key risks for executives, sourcing options for procurement teams, and logistics route adjustments for on-the-ground managers.

With platforms like this, management roles in transportation, storage and distribution (all with a 38.1% exposure score by 2032) could see significant change. The need for human intuition will remain, but AI could streamline data analysis and scenario planning. Even highly physical roles that in general will see relatively lower disruption from generative AI, such as stockers and order fillers (26.4% exposure scores), might see their work augmented by AI-optimized inventory management and warehouse navigation.

[The prescriptive supply chain oracle's strength lies in continuous learning and adaptability. As global situations evolve, the platform evolves with them.](#)



Platform builder


An adaptive learning engine

This AI-native startup would seize on the opportunities created by generative AI's ability to take on at least one-third of the tasks within an educational administration or teaching role—and challenge the one-size-fits-all model of education. In our study, education administrators will see a 36.3% exposure score by 2032, and teachers up to 30.8%.

[Forget standardized tests, and rigid lesson plans—these tools would become obsolete. Instead, the platform would begin with an interactive assessment, uncovering a student's knowledge gaps, learning style and interests.](#) From there, the AI would craft a personalized learning path, with concepts introduced in multiple formats—videos, text, interactive simulations—and adapt in real-time based on how the student responds.

Struggling with a math problem? The AI might offer a visual explanation or break the concept into smaller steps. Excelling in a subject? It might introduce advanced material to keep the student challenged.

But it's not just content; the AI analyzes student engagement, adjusting the pace and even the tone of lessons to optimize learning. The goal is to make education truly student-centric, fostering a love of learning that lasts a lifetime. This type of tool would also find its way into the business world, acting as a personal tutor for employees as continuous learning becomes a critical skill for the workforce in the AI age.



Game-changer

An on-demand sustainable fashion house

This AI-native business would disrupt the wasteful model of fast fashion by combining generative AI with sustainable, made-to-order production. AI-powered design assistants would help customers articulate their style preferences and body measurements and suggest initial design variations, tweaking them in real-time based on customer feedback. Photorealistic visualizations would allow customers to “try on” the garment virtually, seeing it on their own bodies under different lighting conditions.

Once design is finalized, the platform would optimize clothing patterns for minimal waste and connect directly to an on-demand manufacturing facility. Individually made garments would eliminate the overproduction that plagues the industry, and instead, sustainable fabrics, supply chain transparency and design elements that prioritize longevity would be emphasized.

As AI continues to evolve, it's feasible that tasks traditionally performed by fashion designers (43.5% exposure score by 2032) and models (41.6% exposure score) could be streamlined by a platform like this. [Rather than just selling clothes, the platform would offer a radically personalized and ethical fashion experience.](#)



Tech lessons from AI-native businesses

Get ready for the new AI tech stack

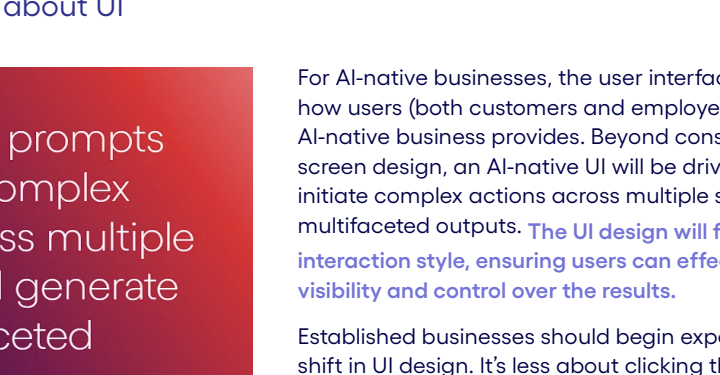
To deliver on these new capabilities, AI-native businesses will build a technology stack that, in many ways, is unrecognizable to businesses operating today. This is because—[unlike traditional businesses that typically see new technologies as a way to do existing things better—AI-native businesses see technology, particularly generative AI, as a way to do things that have never been done before](#).

From an AI-native business perspective, the tech stack is not a static thing that's occasionally tweaked or updated. It's a living, continuously evolving entity, built for a world where agility and continuous innovation are key to a business's survival.

By pulling back the curtain on what an AI-native tech stack might look like, traditional businesses can begin thinking about the changes they need to make to bring generative AI capabilities into their own workplace and, ultimately, their business and operating models.

Application layer

A new way to think about UI



Data layer

Access, connections and creativity

Infrastructure layer

Built for speed and scale

MLOps layer

Enables evolution, not stagnation

AI model layer

An ecosystem of AI models and tools

1. The application layer

A new way to think about UI

Well-crafted prompts will initiate complex actions across multiple systems and generate rich, multifaceted outputs.

For AI-native businesses, the user interface (UI) is a fundamental rethink of how users (both customers and employees) interact with the services the AI-native business provides. Beyond considering button placement and screen design, an AI-native UI will be driven by well-crafted prompts that initiate complex actions across multiple systems or generate rich, multifaceted outputs. [The UI design will facilitate a new human-to-AI interaction style, ensuring users can effectively guide AI agents and retain visibility and control over the results.](#)

Established businesses should begin experimenting with this paradigm shift in UI design. It's less about clicking through menus and more about conversational interaction to achieve goals. To make this mindset shift, businesses can start by asking themselves:

- Where in our existing workflows could a prompt-based UI streamline the experience for customers and employees?
- How do we design a prompt-based interaction that is human-oriented and engaging?

- How can we design a UI that balances the power of generative AI with effective user control and clear feedback on outputs?
- In what ways could a conversational AI interface, directly linked to actions, fundamentally change how users interact with our business and the services we provide?

2. The AI model layer

An ecosystem of AI models and tools

AI-native businesses won't think in terms of a single, monolithic AI model. Instead, they will strategically employ a suite of models, each specialized for a variety of tasks. Such tasks could include text, image and process generation, as well as structured data modeling, prediction and decision optimization, uncertainty modeling, explainability and code optimization.

A diffusion model, for instance, might create stunning product visuals, while a large language model (LLM) would generate tailored marketing messages. Additionally, traditional machine learning (ML) models will be integrated intelligently to address specific problems and leverage the unique advantages of different AI approaches.

Similarly, established businesses should avoid viewing an AI model as a one-size-fits-all solution. It will be crucial to understand the strengths of each model type and select the right tool for the job. Businesses should invest in the expertise needed to determine what purpose these models genuinely serve and how they can be best deployed.

To enter a multimodel world, businesses should consider the following:

- Are we limiting ourselves by focusing only on the latest, trendiest generative AI model? Which specific problems within our business are best solved by a diverse range of AI techniques?
- Do we have the in-house expertise or partnerships in place to understand the strengths and weaknesses of different models, including traditional machine learning?
- How can we strategically integrate various types of models? How will insights from generative models be combined with outputs from other AI or analytics tools?

Avoid viewing generative AI models as a one-size-fits-all solution.

3. The data layer

Access, connections and creativity

Leverage vast quantities of information while ensuring transparency in data acquisition.

AI-native businesses will thrive on information—the wider they cast their net, the more powerful they will become. That's why AI-native businesses will prioritize access to both structured and unstructured data over concerns about data cleanliness. Once they've established clear consent and transparency on how data will be used, they might tap into vast public datasets, third-party resources and customer data.

Additionally, they'll use knowledge graphs (a way of representing data, entities and how they interlink) and vector databases (which are optimized to work with LLMs for fast, easy searches and data retrieval) to grasp complex relationships between data points.

For established businesses, it will be crucial to leverage vast quantities of information, both structured and unstructured. Regardless of how they do this, ensuring transparency in data acquisition will be critical.

However, acquiring massive datasets can be challenging. This is where businesses can get creative. They can use techniques like synthetic data generation, where realistic but artificial data sets are created, often using generative AI itself. They could also use semi-supervised learning, where models learn from a combination of labeled and unlabeled data. By prioritizing access to diverse information and employing these techniques, businesses will empower their generative models to perform at their best.

To improve their data IQ, established businesses should consider:

- Beyond data, are we fully utilizing the information and knowledge embedded within our organization? How can we improve knowledge sharing and information flow across departments?
- What internal data do we possess that no competitor could replicate, and how could that data power our AI models?
- Where do we face data limitations, and could synthetic data generation or semi-supervised learning techniques be viable solutions?

4. The infrastructure layer

Built for speed and scale

Generative AI models can be extraordinarily demanding from a computational power perspective, especially during training and fine-tuning. AI-native businesses will understand that their infrastructure needs to be both scalable and cost-effective, often opting for a hybrid cloud approach. This allows them to ramp up resources for intensive tasks and scale down when usage ebbs, achieving an optimal balance between cost control and performance. Hardware acceleration, utilizing specialized graphics processing units (GPUs) or tensor processing units (TPUs), will be essential.

Established businesses will need to proactively strategize for this level of flexibility and agility. By considering infrastructure implications well before scaling AI projects, businesses can prevent cost overruns and ensure smooth operations.

To make this infrastructure shift, businesses should ask themselves:

- Have we engaged IT teams in proactively forecasting the computational demands of future AI projects, avoiding cost surprises and performance issues?
- Could we benefit from a hybrid cloud approach for our specific AI use cases? When should we prioritize on-demand resources versus reserved instances?
- Does our hardware acceleration strategy (GPUs/TPUs) align with the types of AI models we will be depending on (image generation, natural language processing, etc.)?

Consider infrastructure performance and costs well before scaling AI projects.

5. The MLOps layer

Evolution, not stagnation

Invest in MLOps to ensure your AI systems continuously evolve.

For AI-native businesses, machine learning operations (MLOps) is the cornerstone of maintaining a competitive edge. It isn't a luxury, but rather the framework that allows these businesses to stay agile and responsive. Robust MLOps pipelines enable experimentation, fast deployment of new and improved models, and comprehensive production monitoring. These pipelines will ensure that AI-powered workforces are constantly learning, automatically retraining as new data becomes available, and seamlessly updating to optimize performance.

Established businesses need to shift their thinking toward AI as a dynamic system rather than a one-time project. By investing in MLOps, they can automate processes and shorten the critical feedback loop between real-world results and model improvements. Doing so will ensure their AI systems are continuously evolving to match the dynamic market landscape.

To make the MLOps shift, businesses should consider:

- How can we shorten the time between an idea for model improvement and deploying it in production?
- What are the key metrics to monitor in production, not just for model accuracy, but for real-world business impact?
- Does our company culture empower data scientists and engineers to experiment rapidly, even if there are occasional failures along the way?

A mindset shift

Ultimately, AI-native businesses will instinctively grasp an entirely new way of developing, evaluating, understanding and deploying AI systems. They won't approach AI solutions with the mindset of building rigid, deterministic machines designed for perfectly replicable results. Instead, they will view their AI agents as evolving collaborators with the potential to be far more than mere tools. (For more on this topic, see Cognizant VP of AI Research Risto Miikkulainen's [recently published article](#), "Generative AI: an AI paradigm shift in the making?")

This approach will shape their strategies in fundamental ways. AI-native businesses will prioritize building systems capable of continuous learning, understanding that "training" is never truly over. And they will design interactions that leverage the power of prompts and iterative feedback, guiding their AI agents with the finesse one might use when working with a highly skilled human specialist.

Similarly, established businesses should look to truly harness the power of generative AI rather than shoehorning AI into existing processes or evaluating it through the lens of old metrics. Doing so will require technology teams to think in a completely new way: a relentless willingness to experiment, shift away from a need for total explainability and accept that trust in the AI agent will be earned through rigorous observation rather than upfront programming.

1: Build: The path of ultimate control

The build-it-yourself approach leads to a bespoke AI solution that aligns perfectly with business needs.

It also comes with high costs: attracting top talent, building robust infrastructure and slower time-to-market.

This approach makes sense when AI is critical to competitive advantage or when highly sensitive data mandates in-house protections.

2: Partner: Access expertise and speed

Partnering with AI-native businesses is a faster route to cutting-edge algorithms, robust datasets and expert support.

While integration and vendor lock-in must be considered, this strategy shines when speed is critical and AI expertise is lacking.

AI-native businesses themselves rely on partnerships for specialized functions so they can focus on their core innovations.

3: Utilize: Rapid prototyping and agility

Utilizing pre-built AI solutions via APIs and services is ideal for businesses looking to supercharge development, rapidly prototype and add AI capabilities without reinventing the wheel.

Speed and cost-effectiveness often outweigh the trade-off of limited customization, especially for testing new markets or product enhancements.

4: The hybrid approach: A little bit of everything

Most established businesses will find success by strategically combining elements of all three approaches.

This might mean building bespoke models for core differentiators and relying on pre-built solutions for common tasks.

The ability to pivot as needs evolve is a hallmark of the AI-native business mindset.

Build, partner, utilize

Strategic considerations for businesses in an AI-native world

In addition to thinking like—and ultimately acting like—an AI-native business, established companies will want to partner with AI-native businesses. While a relatively smaller number of traditional businesses may opt to become disruptors by building out their AI capabilities on their own—or spinning off an AI-native division of their own business—a larger number will look to partnerships to bring the fresh perspective of the AI-native in-house.

The great majority will likely take a hybrid approach: partnering with AI-native businesses, using APIs to bring new AI services into their own environment and discerning where they absolutely must build AI capabilities themselves.

In all cases, established businesses can harness AI to transform existing processes and build innovative new offerings—and find ways to integrate the strengths of both traditional and AI-native models to secure their success.

The coming of an AI-native world

AI-native businesses aren't here yet. But there's no doubt they're taking shape as we speak. When they do emerge, they will quickly make our everyday actions look and feel very different, from how we prompt for information and ask for action, to what we get in return. In the background, it will be armies of AI agents, orchestrating and collaborating to get work done more efficiently, productively and unexpectedly than anything seen today.

The thing about being a native is you've got to be born into it—no legacy technology, deep-benched workforces or entrenched mindsets to modernize, reskill or evolve. But what existing businesses can do is learn from these emerging businesses to insert generative AI into their own operations and, ultimately, do business in a whole new way.

By doing so, existing businesses won't just be part of a "remember when" exercise. They'll be a full-fledged part of what is yet to come as AI-native businesses and the organizations of today, together, imagine and create our next new normal.

About the author

Duncan Roberts

Associate Director, Cognizant Research

Duncan Roberts is a Associate Director in Cognizant Research. He joined the company in 2019 as a digital strategy and transformation consultant in industries ranging from satellite communications to educational assessment. He has advised clients on utilizing technology to meet strategic objectives and discover the art of the possible through innovation.

Before Cognizant, Duncan worked for one of the largest publishing houses in Europe, playing a leading role in the digital publishing revolution, helping transform their operations end-to-end and launching new innovative products. He holds a Master's in Philosophy and Classics from the University of St. Andrews.

in ✉



Acknowledgments

The author would like to thank the following for their contributions to this report:

From Cognizant

- Catrinel Bartolomeu, Director of Storytelling and Content
- Mary Brandel, Editor
- Babak Hadjat, Chief Technology Officer for Artificial Intelligence
- Lynne LaCascia, Head of Brand and Creative Strategy, Content, Thought Leadership & Research
- Risto Miikkulainen, VP, AI Research
- Ollie O'Donoghue, Senior Director, Cognizant Research
- Naveen Sharma, Global Practice Head, AIA Practice
- Mike Turner, VP, Software and Platform Engineering