

Building trust into conversational AI solutions

Conversational AI has the potential to transform customer and employee experiences, supporting business growth. But the technology is also sounding a clarion call for a more trustworthy approach. Humans need to listen—and respond.

The power of language

Language is one of the most fundamental of human traits.

From birth, we learn to use language—conveyed in sounds, gestures, and words—to communicate with the world around us. Later, as we learn to read and write, we acquire the ability to use increasingly sophisticated language and a range of symbols and media to interact with other people. And as adults, so much of what we do in our professional, social, and personal lives is communicated or achieved by expressing ourselves through language.


Our relationship with the inanimate world has changed, especially when it comes to the tools and machines, we use day to day. Here, natural language has historically played a far less significant role. For most of human history, tools were simply manual or mechanical. And even as the digital age developed, our interaction with machines has overwhelmingly relied on inflexible interfaces and/or learning to code in the machine’s own “language”.



The advent of natural language processing (NLP) and conversational artificial intelligence (AI) fundamentally changed this dynamic. At last, here was a means for humans to interact with digital systems, through speech or text, in a way that was natural, intuitive, and required no prior “learning” of tools and commands.

As NLP technology has advanced, the early and somewhat underwhelming experiments with narrow and clunky chatbots are giving way to smarter solutions that support more human-like relationships. And more recently, NLP has been supplemented with a new generation of intelligent conversational capabilities, such as the ability to recognize, interpret, process, and simulate human feelings and emotions (affective computing), understand visual inputs (e.g. facial recognition) and process/generate text in ways that were difficult to anticipate just a few years ago (e.g. GPT-3, BERT, etc.).

These new technologies can support even more human-like relationships as they strive to build rapport, show empathy, drive collaboration, and deliver mutually beneficial outcomes for companies and consumers. As we look to the future, we can consider their importance in a metaverse where the conversational sophistication of digital avatars shapes the immersive experience of the customer.



New possibilities, new risks

Unsurprisingly, these new advantages are helping to drive significant growth in the market for conversational AI solutions. By some estimates, the conversational AI industry will be worth \$13.9 billion by 2025.¹ As the market grows, technologies will continue to improve, increasingly blurring the lines between human and virtual interactions.

This is fundamentally changing the human-machine relationship. As machines take on more human qualities, a wide variety of ethical considerations comes into play—including stereotyping, data privacy, behavioral inference, manipulative behavior, and the exploitation of vulnerable populations. These issues can be extremely nuanced and context-specific, which is why accelerated use of conversational AI needs to be matched with an equally accelerated and comprehensive consideration of ethics, effective solutions, and continuous monitoring.



In response to these changes, new AI regulations are emerging that could place far greater requirements on developers of conversational AI. In Europe, for example, a proposal for a comprehensive legal framework on the development and use of AI is under review (see appendix). If approved, the AI Regulation could impact conversational AI systems in ways that range from transparency stipulations to high-risk compliance requirements to the outright prohibition of particular AI practices. Similar plans for regulation of AI are also beginning to emerge in other countries, such as China, the United Kingdom, the United States and Brazil.

In this context, the central question facing business leaders is: **How can we capture the opportunities of advancing conversational AI, while simultaneously addressing the ethics risks and regulatory requirements?**

Conversational AI is buzzing

Businesses in a wide variety of industries are deploying conversational AI solutions to become more attentive and responsive to customers, enhance efficiency, and support growth.

For example, companies such as UBS, BMW, Southern Health Society and Noel Leeming's Stores have partnered with Uneeq, which has developed "digital human" technology, capable of replicating human emotions, gestures and visual cues, in support of real-time conversations.²

Meanwhile, Entropik is working with a number of leading consumer goods companies to capture facial expressions, eye movements, voice tonality and even EEG brain activity readings to better understand user behavior.³

United States insurer MetLife has implemented Cogito's emotion AI coaching solution in ten of its call centers to provide real-time guidance to agents while speaking to customers. Results include a 14-point improvement in its net promoter score (NPS), a widely used market research metric that asks respondents to rate the likelihood of them recommending a company, product, or a service. The company has also increased its "perfect call" scores by 5%, achieved 6.3% greater issue resolution, and a 17% reduction in call handling time.⁴

And the barriers to implementing AI solutions continue to drop. Cloud platforms have made the technology increasingly easy to implement, meaning companies can now spin up new conversational AI solutions with remarkable speed. Accenture helped one European Rail company create a multilingual chatbot to provide critical travel guidance to support the huge influx of refugees resulting from the 2022 Ukraine conflict. The entire solution was designed, delivered, and went live in less than two weeks.



The path to better informed...

... more responsible decision-making.

Recent responsible AI initiatives demonstrate business interest in finding answers to these challenges. The Business Roundtable, for example, has made a series of policy recommendations for governments, as well as setting out various principles that companies should consider when embarking on AI development.⁵

However, when it comes to developing conversational AI solutions on the ground, companies can struggle to translate principles into practice—especially in developing techniques and

processes that ensure ethical standards are adhered to.

Companies can also find it difficult to fully comprehend the unintended consequences of design decisions, given the myriad variables and tension points that need to be considered. In the absence of industry standards and clear regulatory guidance, designers that want to take advantage of conversational AI opportunities have to date lacked a systematic way to identify and address ethical risks.



To help fill that void, Accenture has developed a practical approach to help designers and leaders think through some of the ethical implications and help inform their decisions as they develop and deploy conversational AI tools.

Our methodology considers the intricacies of technology development and human rights in tandem, in the context of conversational AI. We focus in particular on individuals' rights to equality, privacy, and freedom of opinion, choice, and thought.

Starting with Accenture's definition of AI as a technology or system that senses, comprehends acts, and learns, we frame our approach at a high level with three questions:

1. Does the conversational AI have a **human-looking** avatar that might embed stereotypes?
2. Does it set out to **understand the human** user?
3. Does it **behave like a human** in a way that changes the relationship with the end user?

While it is not enough on its own to ensure conversational AI has been implemented in a responsible and trustworthy manner, it can be an invaluable tool in identifying and addressing the additional implications of design decisions, as part of a broader organizational Responsible AI framework.


Looking human

The relationship between the “humanness” of a machine’s appearance and the extent of its appeal to human users is not strictly linear.

Research has famously shown that user affinity plummets into an “uncanny valley” when machines that look nearly *but not entirely* lifelike induce feelings of unease and repulsion.⁶ Nonetheless, anthropomorphism can be a useful tool in building empathy and trust. In the right context, mimicking human features and characteristics in the visual manifestation of a conversational AI can be a valuable way of creating rapport with users.

That said, great care must be taken to intentionally consider how **stereotypes and discrimination** can manifest in an avatar that looks like a human.

For example, consider that a conversational AI assistant is in a fundamentally subservient position to its human interlocutor. It’s there to serve, to respond to requests on demand, 24 hours a day, 7 days a week. Why, then, do so many virtual assistants sound like or look like women? Historically, developers have offered the justification that studies show people prefer interacting with a female voice, although the validity of these findings is disputed.⁷ Either way, choosing a female to embody an “assistant” risk embedding and reinforcing outdated gender stereotypes.



The way an assistant responds to **sexualized speech or inappropriate questions** is another ethical minefield. Research summarized in a recent UNESCO report highlighted how female-sounding voice assistants often respond to abusive language with playful evasion at best or flirtation at worst.⁸

Similar issues can arise with body shape, accent, pitch, and tone of voice. In fact, just about any decision a designer makes about the look or voice of a conversational AI assistant opens up questions of stereotyping and unconscious bias.

Some designers are responding by simply avoiding human-looking avatars at all, even in cartoon form.

Similarly, Google has de-gendered its voice assistant options, referring to them simply by different colors. Accenture, too, has been developing non-gendered voices for digital assistants, including Sam, the world's first non-binary voice solution.⁹

The point is not that conversational AI developers must choose non-gendered voices and avatars, but rather to be aware that there's a whole package of issues to consider. No decision is entirely neutral, and each choice must be thought through and weighed up on its own merits.

Looking human

Key questions to consider:

- **What visual identity and personality are we choosing for our AI assistant and why?**
- **How does this identity support the goals for the interaction?**
- **What accent, pitch, pace, and tone of voice is appropriate?**
- **What unconscious biases might affect decisions about visual appearance?**
- **What is inappropriate or abusive language and how should the assistant respond?**
- **How well does the assistant align with brand engagement objectives and create user stickiness?**
- **Has the assistant been integrated into existing processes in a way that considers employees?**



Understanding the human

Data is fundamental to conversational AI. In order to communicate in a more natural and fluid way, a digital assistant needs to be able to understand and respond to its human interlocutor in real time.

To support this, emerging technologies, such as emotion AI or affective computing, can attempt to infer a human user's emotional state through facial recognition, voice tone, text, and other physiological metrics. Understanding a person's mood in this way can be hugely beneficial in determining the right way to engage. However, this mode of learning raises serious considerations relating to accuracy (is the inference scientifically robust), legality (does it infringe the user's legal rights), and ethics (is it the right thing to do).



User awareness and control are also important issues. Today's conversational AI systems can infer a huge amount about an individual extremely quickly, without that person realizing the depth of that discovery. This is especially true as passive "always-on" data-gathering sensors in smart devices, wearables, and smartphones become more prevalent. In addition, as human users become more familiar with the technology, their willingness to voluntarily share information grows as they increasingly come to think of the system as a companion.

This increases the need for safeguards to protect the **right to privacy** and guard against intrusive profiling.

Specifically, developers should keep a number of key questions front of mind—how much data we are collecting, what purpose is it for, is it more than we need, where are we storing it, for how long, who will have access to it, and so on.

Advancing computational technologies are one potential solution here. In particular, the advent of edge computing—meaning the advanced processing of data on individual smartphones and other devices rather than cloud—allows highly advanced analytics, including machine learning, to be performed on a person's private data without it ever leaving their device. For example, federated

learning carries out processing and machine learning directly on edge devices to maintain data privacy, returning only high-level insights to the centralized cloud. Similarly, Snips, a voice AI platform acquired by Sonos, has the potential to carry out all of its data processing on-device.¹⁰ Indeed, there's now a growing array of technologies, including multiparty computation and homomorphic encryption, that enable privacy-preserving data sharing.

It's also now possible to remove sensitive aspects of voice data before analysis. For example, researchers from Imperial College London have developed a privacy-preserving layer that "sanitizes" voice inputs before sending to the cloud, reducing

the amount of sensitive emotional data captured from the speaker by around 96 percent. And the open-source Smart2 Speaker Blocker from the University of Mississippi is an open-source offline solution that intelligently filters sensitive conversations, preventing them ever reaching a smart device's microphone.¹¹

Algorithmic bias is another key consideration. Preexisting societal biases and inequalities can be reflected in the data used to train conversational AI algorithms. It's therefore essential to consider the risk of bias, whether you're training your own conversational AI system, or purchasing off-the-shelf software.

Understanding the human

Key questions to consider:

- ➔ **How much data is being collected? Are we inadvertently collecting more than we need? And if so, how are we using it?**
- ➔ **Are we being transparent and have we established clear user consent?**
- ➔ **What steps are being taken to mitigate bias?**
- ➔ **To what extent are we making inferences about a user's emotional state?**
- ➔ **How is privacy and data security ensured? What access controls are we implementing?**

A person is riding a bicycle in a dark space. The background is a wall of glowing blue lights arranged in a grid pattern. The person is silhouetted against the lights, and the bicycle is also silhouetted. The overall mood is futuristic and high-tech.

Behaving like a human

The ability to recognize certain aspects of human behavior and respond in a human-like way is key to a conversational AI assistant's value.

By simulating and stimulating emotions in a human interlocutor, the assistant is better able to engage end users and hold more natural sounding conversations. For example, chatbots have shown promise in helping children with autism disorders and people recovering from trauma.¹²

To quell loneliness and improve health and quality of life, Accenture Song worked with Sweden's largest Energy Supplier Stockholm Exergi, to create a conversational AI called Memory Lane that invites a person to share their stories. The AI understands the correlations between different answers to trigger relevant follow-on questions. Every day, Memory Lane analyzes the previous conversation to get to the heart of the topic and to create a memory graph – a virtual and structured version of a person's memory. To preserve these memories, the technology has a built-in feature to translate them into a physical book and podcast to share.¹³

But there are obvious risks here as well. First and foremost, there's the risk of losing the trust of customers or employees. If, for example, a conversational AI is "too effective" at its task, a person may not realize they're speaking with a machine rather than a real human. Some may be fine with this. Others may feel frustrated, embarrassed, or angry when the penny drops. They may feel as if they've been lied to.

For example, Google's Duplex AI project initially demonstrated an impressive ability for AI-powered voices to autonomously hold extended conversations with real people in real-life situations, booking a table at a restaurant, organizing a trip to a hair salon, and so on. But the fact that it didn't disclose to the human participants that they were speaking to a machine proved controversial. The company has since instituted a policy to always disclose that an AI caller is not a human.¹⁴

Another consideration? The extent to which the anthropomorphic features built into a conversational AI encourage an emotional bond. This sense of companionship can be an explicit selling point of a conversational AI service. Xiaolce, for example, is a service that lets users interact with an AI-powered girlfriend or boyfriend via WeChat.¹⁵

On the flip side, think of the implications for an AI tool designed to market a certain set of products or services. At what point does the human user start thinking of the AI as an objective trusted advisor? When might they start to lose a sense of perspective? Performed at scale, end-user profiling, real-time affective analysis and affective simulation can lead to **"hyper-nudging,"** so-called **"dark patterns"** and biometric psychography¹⁶ which exploit users' emotional and cognitive biases. Once again, regulators are taking action with dark patterns prohibited on the interface of online/internet platforms, under the proposed EU Digital Services Act (see appendix).



In these scenarios, the way information, feedback and choices are presented to the user often determines whether an ethical line is crossed.

To begin addressing these concerns from a design perspective, it can be helpful to think of influence in terms of transparency, agency, and recourse. For example, what often separates persuasion from coercion and other forms of influence is how transparent we are being as designers and what choices are being allowed to the end user.

Finally, companies that deploy conversational AI need to anticipate **bad actors**. Bad actors have already demonstrated how they can exploit weaknesses in conversational AI design. Consider Microsoft's Tay Twitterbot, which the company took offline less than a day after its launch in 2016 upon finding that groups of human users had been "training" it on prejudiced and inappropriate content.¹⁷ More recently, Lee Luda, a South Korean chatbot, was suspended after accusations that it used hate speech in its conversations and was being targeted by manipulative users.¹⁸

Bad actors may also be able to manipulate conversational AI in unexpected ways, such as by supplying input data that has been carefully designed to hack the system or engineer a particular result. This is similar to the way self-driving cars can be fooled by, for example, subtle changes to road signs that make them think a 30-mph speed limit is in fact 80 mph.

Behaving like a human

Key questions to consider:


- **How open and transparent is our use of AI in each interaction?**
- **What level of agency does the user have?**
- **Are we adequately preserving their freedom of opinion, choice, and thought?**
- **What recourse does the user have if they want to make a different choice?**
- **Have we cast a wide net when thinking about the ways bad actors could engineer undesirable results?**

Putting ethics into practice

Who's responsible and how to get started.

Ultimately, when it comes to bringing ethics to bear on conversational AI, companies need to think bottom-up and top-down. Our methodology enables product development experts on the ground to start taking action today, helping to identify and act on tension points, while broadening horizons for leadership teams.

However, to ensure decisions are taken in a way that consistently reflects the ethics and values of the company, C-level executives will need to build out a more complete responsible AI framework. Such a framework should include establishing a governance structure. It should also articulate roles and responsibilities, recognizing that higher level ethical decisions will often require input from many different stakeholders and functions—including the C-suite, product owners, designers, developers, and data scientists.



Think about scale, frequency and impact.

What is the scope of the solution's user base—millions of customers or a few hundred employees? The greater the reach, the greater the chance of even micro-level choice architectures impacting large numbers of people. And the more personalization at scale, the greater the potential for conversational AI to misstep in individual interactions.

Consider the company brand and values.

Conversational AI represents a new way for consumers or employees to interact with the brand. It's much broader than a simple technology implementation. How do these tools support and enhance the brand? Does it represent its values and principles?

Build diverse and multidisciplinary teams.

Conversational AI development requires a range of skills and expertise. Look to bring in conversation designers, data scientists and ethics advisors, as well as those who can provide broader input on questions of diversity and the lived experience of different groups.

Consider data literacy.

Ethical AI is not merely a question for data scientists and developers. The whole business needs to be able to understand and weigh up the risks. For example, recognizing a conversational AI with "80 percent" accuracy will provide a "wrong" result every five attempts. The whole business needs to appreciate the ethical implications of this kind of finding.

Expect the unexpected.

As conversational AI systems become more human-like, human users may feel they can reveal more personal information, about their health, their finances, their emotional state, threats to their wellbeing and so on. What happens if and when a conversational AI system designed for a narrow functional use case starts picking up information about a person being abused or considering self-harm? Remember that natural conversation evolves in inherently unpredictable ways. It's important to consider and plan for as many potential outcomes as possible, however far removed they may seem from the initial use case.

Don't underestimate the power of words.

Conversational AI is unique in that words are the interface. And words are powerful. It's easy to inadvertently exclude users or cause offence by using language that may be perceived as sexist, racist, or otherwise biased, prejudiced, or denigrating. Look to involve as many stakeholders as possible during the design process and make sure the content the system serves up to users has gone through a rigorous inclusivity and accessibility review.

Ask who benefits and who has control.

If the end user gets a direct benefit from a conversational AI experience, they may be more willing to accept trade-offs in other areas (such as giving a company access to their data). But if it's the company that is primarily benefitting, justifying the approach through the ethical framework becomes even more important.

Consider the suitability of the AI solution.

Is the technology scientifically valid? And is it appropriate to the use case? Question marks remain over what can be accurately inferred through affective computing techniques; for example, a smile does not necessarily infer happiness. Furthermore, an AI system with all the bells and whistles may be overkill in many situations. The technology should do only what is required to meet the objectives of the experience being designed.

Don't leave conversational AI unattended.

Recognize that conversational AI systems are living systems that need constant nurturing and refinement. Consider setting up an ongoing capability that will consistently monitor algorithms for bias and ensure the system can adapt to new linguistic trends among different demographics and generations (such as the use of emojis).



Trusted and trustworthy AI



Conversational AI is on the cusp of fundamentally changing the way machines can support and improve human lives.

Its potential to improve customer and employee experiences is profound. As we have seen, the advance of the technology is also sounding a clarion call for a more trustworthy approach.

Humans need to listen and respond.

As noted in Paul Daugherty's book "[Radically Human](#)", trust is set to become a key differentiator for

AI companies, and those that fail to adapt will ultimately be left behind. However, trustworthiness is more than just risk mitigation and regulatory compliance. It is a complex topic incorporating competency, reliability, integrity and empathy, requiring a comprehensive Responsible AI framework. As part of this approach, we must consider how trust is manifested in design decisions, which shape how conversational AI looks, understands and behaves, and the subsequent implications for workers, users and society at large.

Appendix

EU Artificial Intelligence Act

On the 21st April 2021, the European Commission published its proposal for an Artificial Intelligence Act.¹⁹ If approved, it will take a risk-based approach and differentiate between uses of AI according to four different risk categories:

- **Unacceptable risk.** AI systems that constitute a clear threat to the safety, livelihoods, or fundamental rights of people. This includes the use of systems with significant potential to manipulate human behaviour and actions, particularly of vulnerable people.
- **High risk.** AI systems are considered high-risk (i) if it concerns (a safety component of) a product already subject to EU safety regulations, or (ii) if designated by the European Commission as high-risk – such as applications that affect critical infrastructure, public safety, medical interventions, law enforcement, employment rights, or essential services.
- **Limited risk.** AI systems with specific transparency obligations, such as chatbots and deepfakes. Users must be made aware that they are interacting with a machine so that they can make informed choices or decide to step back.
- **Minimal risk.** Applications that represent minimal or no risk for citizens' rights or safety remain unregulated.

Digital Services Act

On 23rd April, 2022, the EU reached political agreement regarding the EU's Regulation on a Single Market for Digital Services, or the the Digital Services Act (DSA). This regulation will include the prohibition of dark patterns. "Providers of online platforms will be required not to design, organise or operate their online interfaces in a way that deceives, manipulates or otherwise materially distorts or impairs the ability of users of their services to make free and informed decisions."

Further reading

For more details on the topics covered in this report, please read the following:

[From AI compliance to competitive advantage](#), by Ray Eitel-Porter, Accenture's Global Lead for Responsible AI, et al

[The EU's AI Act: an initial assessment](#) by Ray Eitel-Porter, Accenture's Global Lead for Responsible AI.

[Getting Emotional: How Platforms, Technology, and Communications companies can build a responsible future for Emotional AI](#), by Robin Murdoch, Accenture's Global Industry Managing Director, Software & Platforms, et al

[Scaling with AI ethics in mind](#), by Fernando Lucini, Accenture's Global Lead Data Science & ML Engineering – Applied Intelligence

[Maximize collaboration through secure data sharing](#), by Teresa Tung, Accenture's Cloud First Chief Technologist, et al

For more on Accenture's interdisciplinary approach to AI ethics and governance, please click [here](#). And for insights into designing a responsible AI framework, please read Accenture's point of view [here](#).



Authors



Laetitia Cailleateau

Lead – Data & AI, Europe



Patrick Connolly

Research Manager AI Research

References

¹“Conversational AI Market by Component (Platform and Services), Type (IVA and Chatbots), Technology (ML and Deep Learning, NLP, and ASR), Application, Deployment Mode (Cloud and On-premises), Vertical, and Region - Global Forecast to 2025.” Markets and Markets, June 2020. Retrieved from: <https://www.researchandmarkets.com/reports/5136158/conversational-ai-market-by-component-platform>

² <https://digitalhumans.com/what-are-digital-humans/>

³ <https://entropiktech.com/case-studies/>

⁴ <https://cogitocorp.com/insurance/>

⁵“Business Roundtable Artificial Intelligence <https://www.businessroundtable.org/policy-perspectives/technology/ai#:~:text=Business%20Roundtable%20also%20recognizes%20that,be%20a%20continually%20evolving%20journey>

⁶“Experiments reveal why human-like robots elicit uncanny feelings.” Science Daily. September 10, 2020. <https://www.sciencedaily.com/releases/2020/09/200910110857.htm>

⁷“I’d Blush If I Could.” UNESCO, 2019, <https://en.unesco.org/ld-blush-if-i-could>

⁸ Ibid.

⁹“Accenture and CereProc Introduce and Open Source the World’s First Comprehensive Non-Binary Voice Solution.” Accenture, December 16, 2020. <https://newsroom.accenture.com/news/accenture-and-cereproc-introduce-and-open-source-the-worlds-first-comprehensive-non-binary-voice-solution.htm>

¹⁰ Kinsella, Brett. “Sonos Acquires Snips for \$37.5 Million.” November 21, 2019. <https://voicebot.ai/2019/11/21/sonos-acquires-snips-for-37-5-million/>

¹¹“The Smart2 Speaker Blocker: An Open-Source Privacy Filter for Connected Home Speakers”, Ilesanmi Olade, Christopher Champion, Haining Liang, Charles Fleming, 15 Jan 2019. <https://arxiv.org/abs/1901.04879>

¹²“Chatbot apps for communication and social interaction therapy.” CSIRO. Last updated: November 8, 2021. <https://www.csiro.au/en/research/technology-space/ai/chatbots-for-therapy>

¹³ <https://www.accenture.com/ie-en/case-studies/interactive/memory-lane-conversational-ai-solution>

¹⁴“Google’s creepy AI phone call feature will disclose it’s a robot, after backlash.” Mashable, May 10 2018. <https://mashable.com/article/google-duplex-disclosures-robot>

¹⁵“Always there: the AI chatbot comforting China’s lonely millions.” France24. August 24, 2021. <https://www.france24.com/en/live-news/20210824-always-there-the-ai-chatbot-comforting-china-s-lonely-millions>

¹⁶“Watching Androids Dream of Electric Sheep: Immersive Technology, Biometric Psychography, and the Law.” Vanderbilt Journal of Entertainment and Technology Law. 2021. <https://scholarship.law.vanderbilt.edu/jetlaw/vol23/iss1/1/>

¹⁷“I’d Blush If I Could.” UNESCO, 2019, <https://en.unesco.org/ld-blush-if-i-could>

¹⁸“South Korean AI chatbot pulled from Facebook after hate speech towards minorities.” The Guardian. January 14, 2021. <https://www.theguardian.com/world/2021/jan/14/time-to-properly-socialise-hate-speech-ai-chatbot-pulled-from-facebook>

¹⁹“Proposal for a Regulation laying down harmonised rules on artificial intelligence.” European Commission. April 21, 2021. <https://digital-strategy.ec.europa.eu/en/library/proposal-regulation-laying-down-harmonised-rules-artificial-intelligence>

About Accenture

Accenture is a global professional services company with leading capabilities in digital, cloud and security. Combining unmatched experience and specialized skills across more than 40 industries, we offer Strategy and Consulting, Song, Technology and Operations services — all powered by the world’s largest network of Advanced Technology and Intelligent Operations centers. Our 699,000 people deliver on the promise of technology and human ingenuity every day, serving clients in more than 120 countries. We embrace the power of change to create value and shared success for our clients, people, shareholders, partners and communities. Visit us at www.accenture.com.

Copyright © 2022 Accenture. All rights reserved. Accenture and its logo are registered trademarks of Accenture.

This content is provided for general information purposes and is not intended to be used in place of consultation with our professional advisors.

This document makes reference to marks owned by third parties. All such third-party marks are the property of their respective owners. No sponsorship, endorsement or approval of this content by the owners of such marks is intended, expressed or implied.