appian

Implementing Private AI: A Practical Guide

Why traditional AI methods are limited—and what to do instead.

The artificial intelligence arms race.

You can't go very far today without seeing a reference to AI. The door has been opened, and AI has hit the mainstream. And with the release of large language models—most notably ChatGPT from OpenAI—the public consciousness has been gripped by the promise of AI to do everything from take productivity to new heights to facilitate potential breakthroughs in medicine.

But while AI has advanced significantly in recent months and years, it isn't new. Businesses have long been investing in AI. Some organizations have even built their entire business models with a data science–driven approach. Whether they are using a third-party public cloud provider or built their own in-house AI models, many organizations have at least started to adopt AI in their operations.

However, organizations face clear challenges to widespread AI adoption: How can they integrate AI into wider workflows and do so at scale?

And how do they keep data private in a world of public AI models like ChatGPT? Already facing pressure to do more with fewer resources and lower budgets, organizations need to apply AI in the most effective way possible or risk falling behind.

AI needs to be integrated smoothly into companies' existing processes and operationalized at scale—and, perhaps above all else, organizational data needs to remain private at all times. Done properly, this will speed up the wheel of revenue, reduce operating costs, and free workers for more productivity.

This eBook will discuss how organizations across the public and private sectors have traditionally integrated enterprise-scale artificial intelligence and the benefits and limitations of these methods. Then, it will offer a better approach that can yield steady, recurring benefits for your organization.

Traditional AI implementations.

Despite how it may feel, AI isn't the new kid on the block—it's just gaining a better name for itself. From predictive modeling to task automation, organizations have long applied AI to a wide range of use cases. In the past, and still today for some, deploying AI traditionally came down to two options: large public cloud providers or in-house models. Both have benefits and drawbacks.

Large public cloud providers.

Large public cloud providers offer a range of tools and resources to support development, deployment, and management of AI-driven applications. These large-scale providers offer three main benefits:

1. **Computing power:** resources like virtual machines, containers, and serverless computing options that can handle intensive workloads like training AI models.
2. **Storage:** extensive data storage capacity, which can be critical for companies scaling to accommodate rapidly expanding data sets.
3. **AI services:** pre-built services and tools to simplify the development, deployment, and management of AI applications, including pre-built AI models, APIs, and tools for data preparation and model training.

However, public cloud providers share one glaring weakness—data privacy. Providing proprietary or sensitive data to these large cloud providers can be a dicey proposition, as it's not uncommon for them to use customer data to train their own algorithms. (It's worth noting this was enough of a concern that OpenAI changed their policy to explicitly state they won't train their models on customer data for API-developed applications.[1] However, they still may use your input into the ChatGPT chatbot for training that model unless chat histories are disabled.[2])

For many organizations, that lack of privacy makes public cloud providers a non-starter. And organizations in industries under strict data privacy compliance laws will be particularly wary of integrating AI into their workflows for fear of a leak or simply the uncertainty that comes with sharing data without clear parameters around how that data will be used.

1. "Addressing Criticism, OpenAI Will No Longer Use Customer Data to Train Its Models By Default," TechCrunch. (March 1, 2023).
2. "Lawyers Breathe a Sigh of Relief: They Can Turn Off Chat History for ChatGPT," Above The Law. (May 2, 2023).

In-house AI.

Many organizations choose to build their own data models and infrastructure in-house. This approach does offer more privacy and greater control over the technology and infrastructure, but it's also expensive and work-intensive. You'll need a team of experts—data scientists to build and test models, data engineers to prepare and clean the data, software engineers to integrate the data into applications, and IT teams to build and maintain the infrastructure and storage required. Plus, don't forget the cost of hiring a security team to ensure your data remains private.

The process of building an in-house model is complex and time-intensive. It typically includes:

- **Preparing data** to train the model. This involves collecting, cleaning, and preprocessing data as well as ensuring the data set meets the bar for being high quality, representative, and unbiased.
- **Extracting features**, or the process of identifying the input data that the model will use to make predictions or decisions.
- **Selecting a model**, which requires data scientists who know how to apply the right data model based on a given data set.
- **Tuning the model** to ensure that it's performant.
- **Training the model on a data set.** This involves machine learning, where a system teaches the AI model to recognize patterns or relationships in the data so it can then make predictions based on the new data.
- **Provisioning the infrastructure** that will enable the AI model to be used in production.
- **Retraining the model** for continuous optimization. Data drift—where the parameters/data set changes over time—can make the model less accurate and require retraining. You may also need to modify models when new data becomes available or changes are made to tasks the AI was designed to achieve.

Building a bespoke model using an in-house team carries the risk of project failure or scope creep. Plus, the organization will face long-term costs to maintain the model and the underlying infrastructure. Finally, integrating an in-house model into existing workflows can bring challenges of its own. Building in-house AI models isn't just a matter of hiring software developers—the complexity and costs depend greatly on the underlying technologies involved. The more complex the IT ecosystem, the harder it will be.

The missing pieces.

We've discussed the benefits and limitations of two common approaches to AI, but there's something else to consider. Whether using a large public cloud provider or building an AI model in house, each individual AI model is still a single-use tool. Yes, it can help solve incredibly complex problems like generating an image or choosing a price point for purchasing a stock, but it's still only one tool used for one problem. In short, AI isn't enough on its own.

Take the example of trying to streamline a billing process. You could easily apply AI to the problem of classifying incoming billing-related documents. But businesses need more than a local task cleanup. How do you go about automating communications? What about routing this information to the right people? What about simplifying data entry once you've extracted data from the invoices? These tasks require other automation technologies like API integrations, RPA bots, and workflow orchestration to truly streamline the full process and route work to the appropriate people or digital workers.

Connecting underlying systems also brings its own challenges. Often, you'll have to pull data from and push data to multiple sources to accomplish a given task. This can require extensive API development work as well as database programming. If any of these underlying systems change their database schemas, there will need to be a corresponding change in the company's IT ecosystem.

Fortunately, there is a simple solution to filling in these gaps.

The best of both worlds.

So, large public cloud providers can make it comparatively easy to use AI-powered services, but fall short in terms of privacy. And in-house AI development gives you that much-needed privacy, but shifts the burden of building and enabling those models to the consumer. Are you caught between two poor options?

Thankfully, no. Process automation platforms that include native AI capabilities can allow you to train your own private models without having to build an expensive data team in-house (or shell out a pretty penny to build and maintain infrastructure). Plus, there are other, deeper benefits—most notably the ability to incorporate AI easily into wider workflows. Let's take a look at some of the benefits an AI-powered process automation platform can give you.

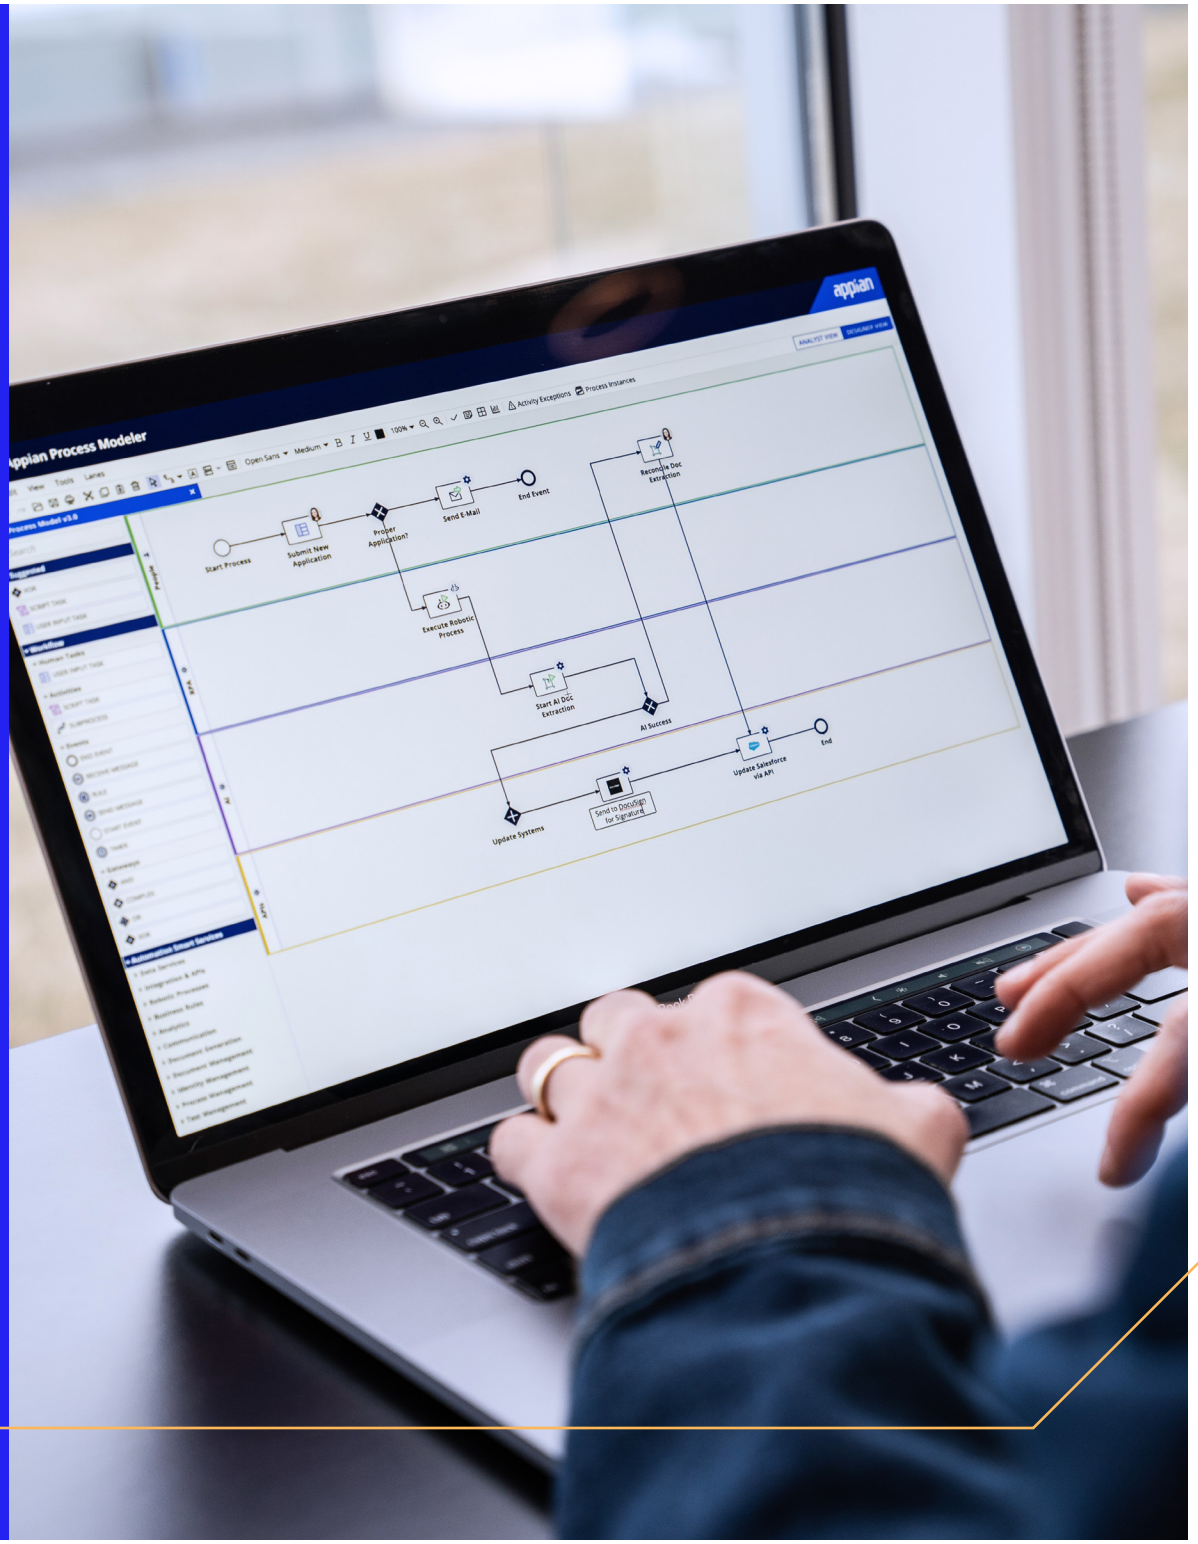

Keep your data private.

As mentioned earlier, when you give your data to large public cloud providers, you sacrifice privacy. These providers often build their business models on the premise of gaining access to data. They use your data to hone their own algorithms. Making matters worse—these algorithms are shared by all their customers, which means your own data could be helping your competition. So, how does a process automation platform with native AI skills address these issues?

First, you're able to create your own private AI models trained on your own data that's never used to optimize anyone else's algorithms. You're creating your own private model, not offering your data to the collective. Each individual model gets trained on your data—and it won't be used to train the cloud provider's larger algorithms. Beyond privacy, this private model generation paradigm is just more practical. Your organization has unique needs. By taking this approach, you tailor your algorithms to you rather than having to use a more generic, publicly available algorithm.

Second, a good platform puts several safeguards in place to secure data. Make sure to have your security team review the security protocols of any platform you're considering so you know that your data's in good hands. One critical feature to look for is private-key encryption, which involves giving you a customer-defined private key that's not available to the platform provider. Because each customer has their own private key, it would be difficult for anyone to decrypt your data in the unlikely event of a data leak. While data breaches will be uncommon for organizations that use preventive, defense-in-depth security features, this backup measure further protects your data from becoming public.

Build AI models at lightning speed.

So, a strong platform can provide privacy. But how hard is it to build a data model on a platform like this? For most in-house private AI implementations not done on a platform, the task is burdensome, with many complex steps that require a team of data experts.

The platform approach eliminates this problem. With a minimal amount of prep-work, you can train and deploy bespoke AI models in minutes. All you need is a sample set of data, then you can simply let the system train on the data.

We'll use the Appian AI Skill Designer as an example here. Let's say you wanted to build a model that classifies incoming emails. The process goes like this:

1. Choose a sample set of emails that are representative of the types of incoming emails you might receive (e.g., invoices, payments, tax information, etc.).
2. Select the Email Classification skill in the Appian AI Skill Designer.
3. Upload the sample emails.

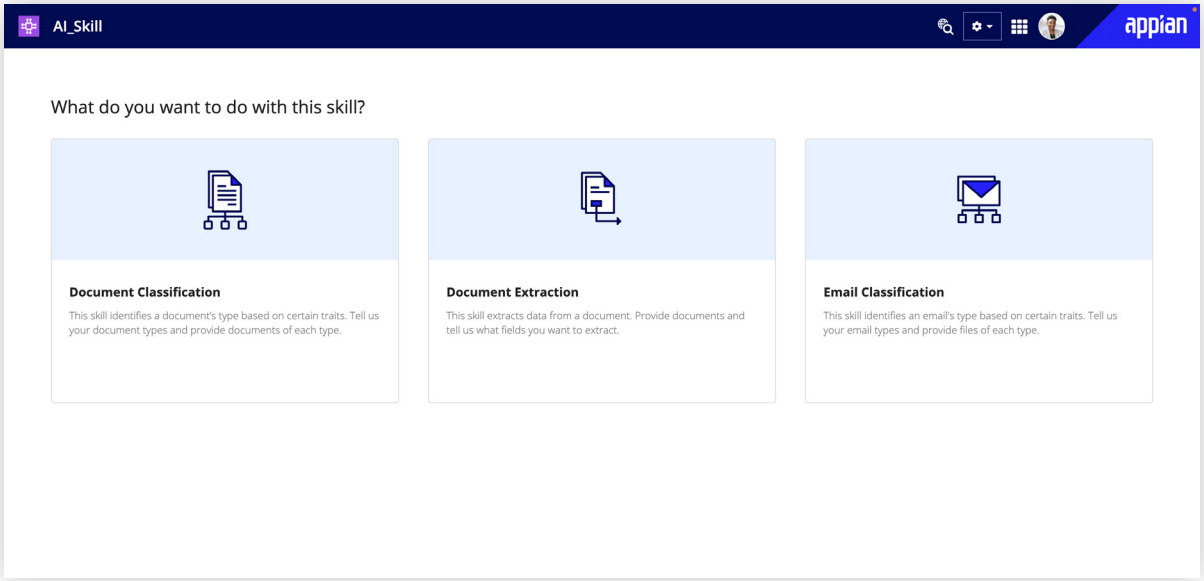

Figure 1: Select from three currently available AI Skills. In this case, you would choose the email classification skill.

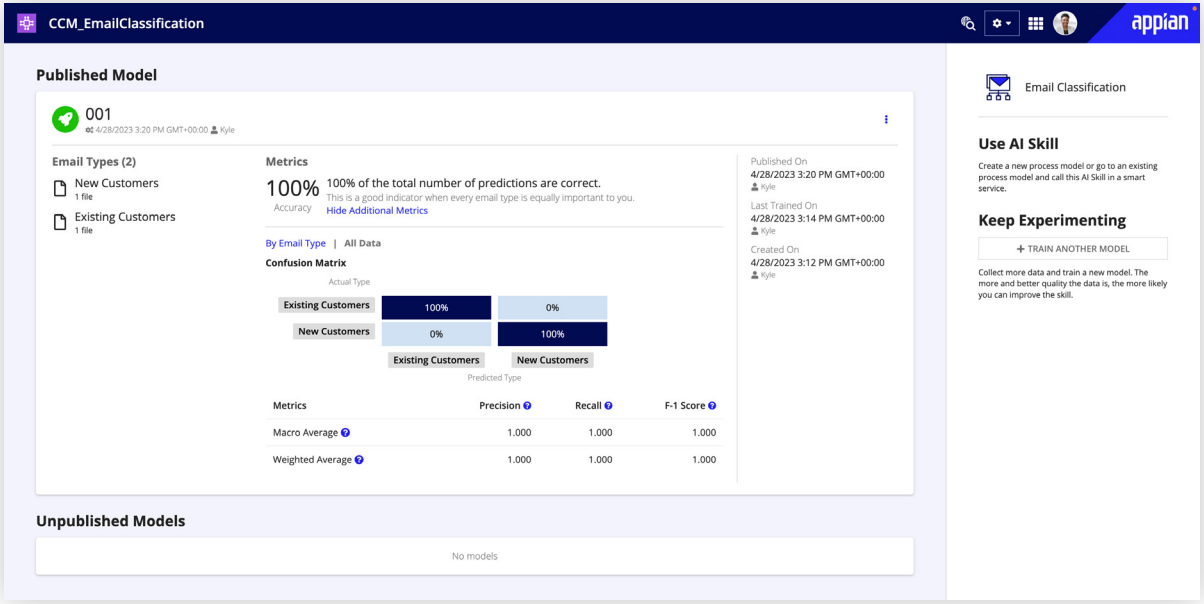

Figure 2: Output metrics from a trained email classification model.

4. Let the model train on the emails—and receive an output in minutes.
5. Review the metrics to ensure the model is ready for publishing.
6. Either decide to release the model or create a new model in minutes to meet your needs.

It really is that simple to get a working AI model that's tailored to your own needs; the process can be done without writing code or developing complex statistical models.

Here's another benefit: Building a model is quick, and retraining a new model is equally fast. AI isn't a set-and-forget technology; you do need to retrain your models over time.

For example, consider data drift. This concept refers to the properties of a data set changing over time, which degrades the model and decreases prediction accuracy. Think about how data drift could affect a document classification task: if the format of an organization's invoice template changes over time, then the model may have trouble recognizing and classifying the document (or extracting proper information from it). A model is only as good as the parameter's it's trained on, and the world is not static. Using native AI in a platform gives you the ability to build a new model in minutes based on updated information and then allows you to push that new model version out fast. The result? Your AI models remain tailored to your specific needs at any moment in time.

Reach the true goal: efficiency.

AI is a revolutionary technology. But let's be clear: simply using AI is not the goal in and of itself. AI is a tool for a greater purpose. The real goal is more efficient processes. And reaching that goal requires automation—whether it's AI, RPA, business rules, or other automation technologies. Automation comes with benefits like enhanced efficiency, increased productivity, and a healthier bottom line for organizations regardless of the economic weather.

In other words, AI is one tool among many. Don't chase it for its own sake. Getting to full process automation requires multiple technologies working together with humans in harmony.

Here's an example to illustrate that point. Let's say you want to tackle optimizing an end-to-end billing process. AI-driven document processing can receive documents, recognize what they are (payments, requests for extensions, etc.), and then extract data from them. If it recognizes a payment, it can then pull data fields from the document, such as the payment amount and the account number. But to credit that information to the person's account, you will need a connection to your main billing system. This requires either API connectors or an RPA bot to input data. And humans need to remain in the loop to review potential errors or discrepancies, deal with issues that require escalation, and potentially communicate with customers directly. And you'll need an orchestration layer to route work between all these technologies and people.

Once again, AI is only one tool—you need a well-stocked tool kit to automate a full process.

A good process automation platform incorporates all these capabilities in one solution. Beyond automation tools like AI, API connectors, and RPA, process automation platforms give you the ability to use business rules and logic to route work appropriately and pass off tasks between humans and digital workers seamlessly. Using a separate AI provider requires a lot of manual integration work—while taking a platform approach allows you to sidestep all this and roll out automations fast.

It takes a platform to launch.

We're biased, but we believe a platform approach is the best way to integrate AI into your business workflows. The Appian AI Skill Designer, specifically, will help you achieve your AI and process goals quickly and headache-free. With it, you can:

- Build and train models rapidly—often in minutes. Upload a sample data set, then review the outcome metrics. From there, you decide whether to publish the model or build a new one. Plus, you can retrain the model at any time without disrupting any live models in production.
- Retain data privacy by creating your own private, tailored models. We never train models on your data. Plus, all data is encrypted via your private key and backed with the strong security processes and certifications of Appian Protect.
- Build AI into a wider digital experience using low-code design. This solves the problems that can come with integrating AI into your wider workflows so you can truly automate end-to-end processes and pass off work between digital workers and humans with ease. Plus, with multiple automation technologies available in addition to AI—such as RPA and API integrations—you always have the right automation tool for the job.

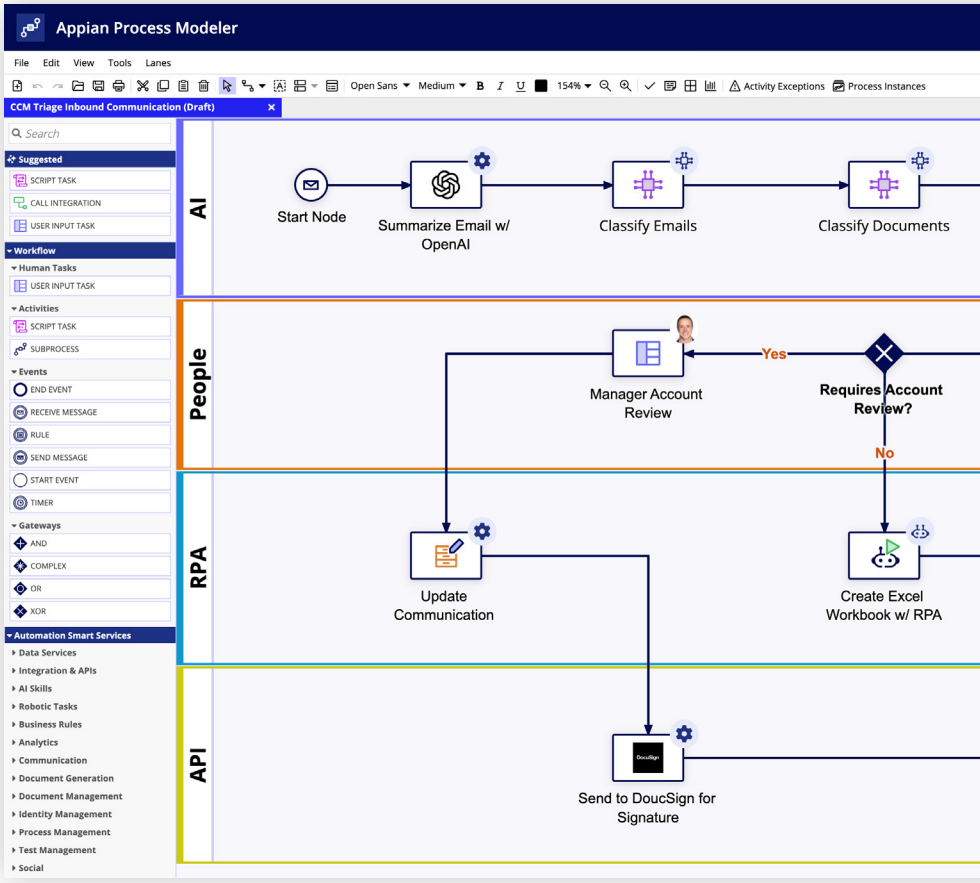

As the AI landscape continues to evolve, businesses face an arms race to integrate AI into their operations. But it's critically important to be smart about how you integrate and operationalize AI. It should be implemented as part of a larger automation strategy that helps truly transform your business. Appian is leading the charge by offering secure, private AI models without the onerous process of traditional in-house development.

Learn more.

You can learn more about the Appian AI Skill Designer by reading the datasheet or watching the demo from Appian World 2023.

If you're curious about our vision for the future of AI and our current roadmap—including how we plan to incorporate generative AI—then check out the Appian World 2023 Product Vision Keynote or visit our AI Vision web page.

And if you're interested in learning more about Appian's approach to AI or other process automation technologies, contact us. We're always happy to talk.

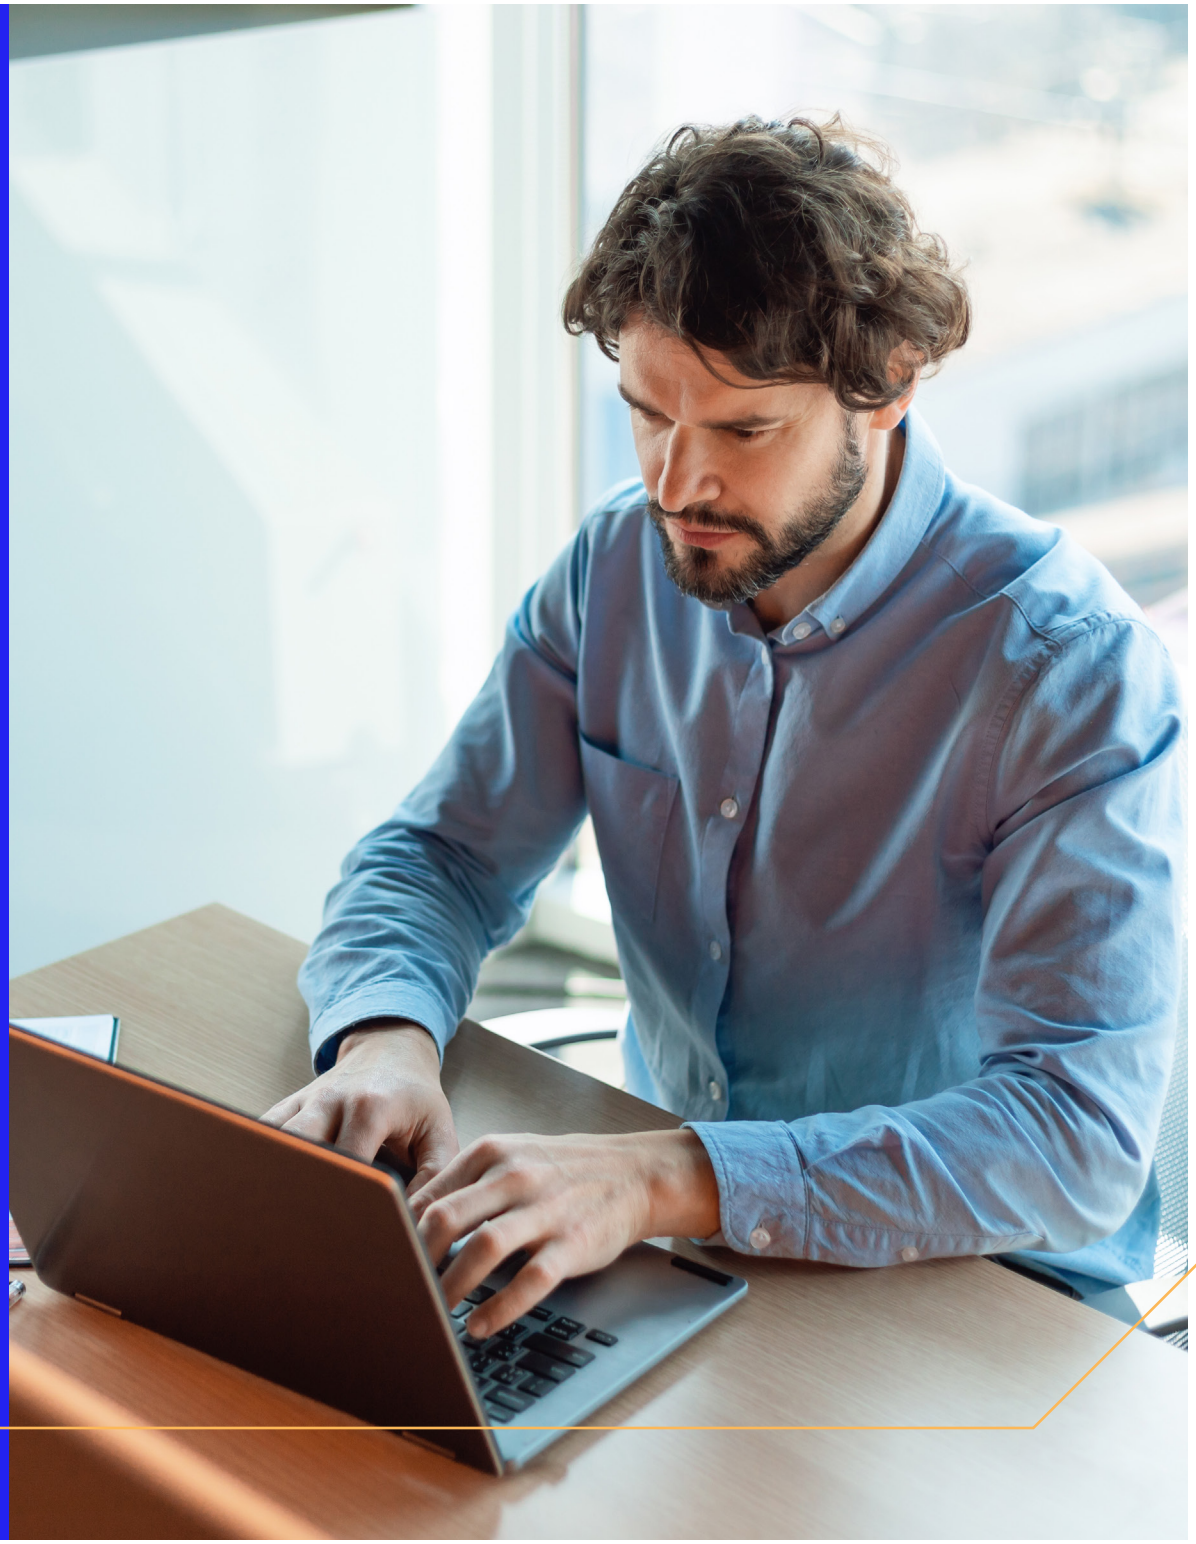

appian

Appian is a software company that automates business processes. The Appian AI-Powered Process Platform includes everything you need to design, automate, and optimize even the most complex processes, from start to finish. The world's most innovative organizations trust Appian to improve their workflows, unify data, and optimize operations—resulting in better growth and superior customer experiences. For more information, visit **appian.com**. [Nasdaq: APPN]

EB-6007426