DATA AUTOMATION GUIDE

# UNLOCKING THE PROMISE OF GENERATIVE AI

Harness the power of genAI without compromising your data

Vendia

# Table of contents

#00

# It goes without saying

that generative AI (genAI) is having a moment. Yet it's the same surplus of shiny genAI tools cropping up across departments and teams that has business and IT leaders so worried. While organizations recognize data security and privacy as top priorities in a genAI world, few are equipped to proceed with confidence.

Among the future unknowns and business uncertainties are emerging questions like:

- How do we keep proprietary data out of online genAI tools?
- How can we explore new use cases without risking privacy and noncompliance?
- What guardrails help prevent confidential data from being exposed within or outside the company?

The answers to these questions are often anything but straightforward, linear, or sureproof—if they were, this guide wouldn't exist! But fortunately, with a few simple measures and the right governance tools in place, enterprises can safely embark on this exciting new frontier while keeping sensitive and proprietary data secure.

# #01

## Generative AI: **Data goes in but never, ever comes out**

Like any disruptive new technology, the use of genAI naturally surfaces concerns around safety, privacy, and security. When used with care and caution, genAI applications can be hugely beneficial across an organization. On the flip side, these tools have the potential to be dangerous and damaging if applied indiscriminately.

Here's why: GenAI applications are powered by machine-learning models, such as large language models (LLMs), which not only understand words but also the relationships between and concepts behind them. Through extensive training on source material—publicly or privately available data or a combination of both—these models then use this information to predict a finite number of answers to a given query.

The trouble is, these models are intrinsically incapable of deleting, redacting, and restricting access to data when answering a query. In other words: once data goes in, it never comes out. Because as powerful as an LLM may be, it simply doesn't know why it knows something or what information shouldn't be revealed.

Worse, once data goes into an LLM, it's there forever and for everyone—accessible to anyone at any time. While efforts to make "selective forgetting" a reality are underway, it'll likely be years, if not decades, before we see such capabilities in production use.

Until then, it's all or nothing: If you somehow let the wrong information get into an internal model, your only solution is to delete the entire thing and start over. And if it gets into a public or shared model? Unfortunately, any slip-ups here become "forever" mistakes.

Free download
Click here