DATA AUTOMATION GUIDE

UNLOCKING THE PROMISE OF GENERATIVE AI

Harness the power of genAI
without compromising your data

Vendia

Table of contents

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**

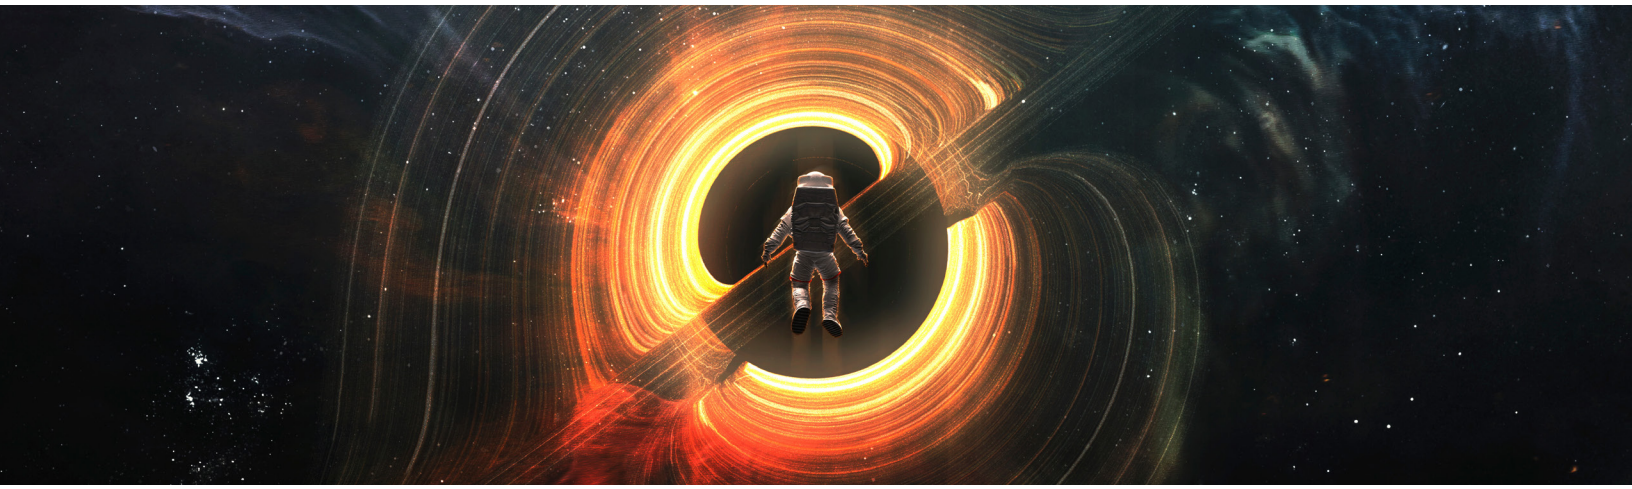

Like any disruptive new technology, the use of genAI naturally surfaces concerns around safety, privacy, and security. When used with care and caution, genAI applications can be hugely beneficial across an organization. On the flip side, these tools have the potential to be dangerous and damaging if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

#02

When confidential data leaks, the consequences pour

The unfortunate reality is there's never a guarantee that company data will never get into the wrong hands or be used without permission. But even more unfortunate are the consequences this now brings in light of genAI—where losses grow bigger and more costly for businesses.

Consider the potential fallout from the following business scenarios:

1. A sensitive report or company spreadsheet leaks online

The ease and speed in which leaked company details can be exploited accelerates via genAI tools—even a small data leak can snowball into significant company and competitive losses.

For example, imagine a financial report or sales account list is accidentally shared online with an LLM, which is then incorporated into its training material and made publicly accessible. Without the company's executive or InfoSec team knowing it, anyone online can now probe for detailed insights regarding the company's sales, customers, or finances.

Even worse, there's nothing company leaders can do to rectify the situation. No alerts or early warning signs of damage exist until the exposure actually occurs. As previously mentioned, not only does an LLM not recognize when it has information it shouldn't, but it also can't be asked to forget any data either.

2. Consumer data is exposed via internal training models

GenAI brings about a more insidious risk of data leakage when it comes to consumer privacy and compliance. This is especially true in highly regulated environments such as the healthcare, pharmaceutical, and financial services industries, where LLMs are often trained on a combination of public and private datasets containing personally identifiable information (PII) and personal healthcare information (PHI).

Let's say a healthcare insurance company builds an online portal tool for patients to locate the right provider service based on symptoms. Doing so would require a training environment to first train the LLM on millions of patient records, along with a production environment to then deploy the model for external use. If PII and PHI haven't been properly redacted from the training model, this information could accidentally leak into the deployed model for future use. Meaning that confidential health details of public figures, family members, and coworkers could then become easily available for anyone to find.

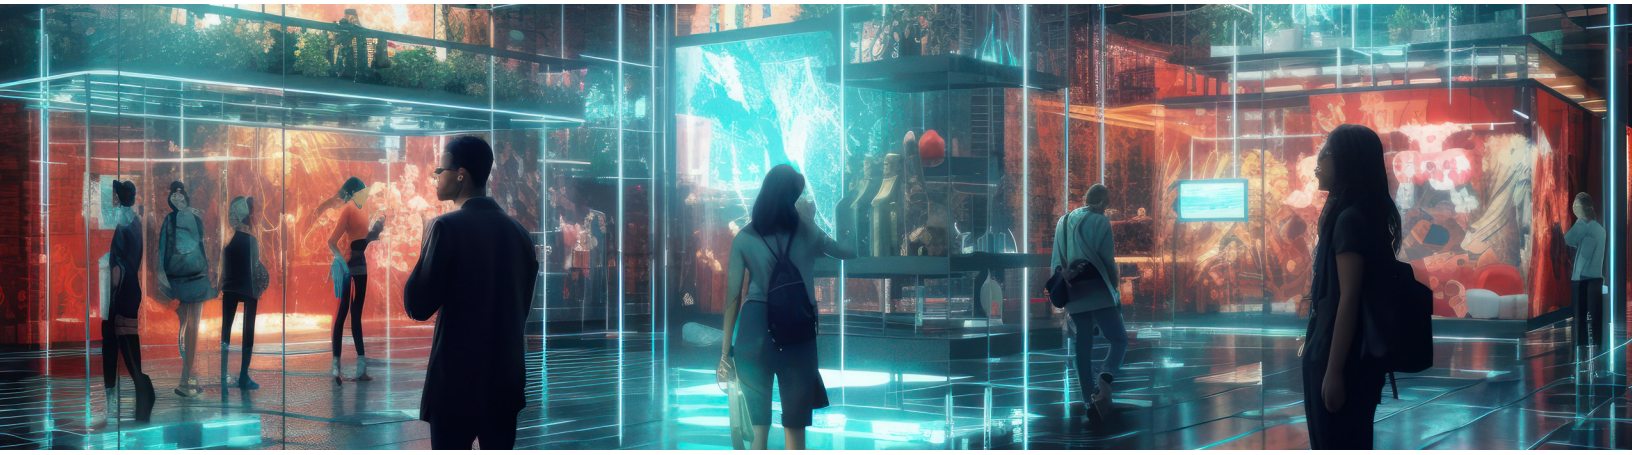

3. Decisions become influenced by false assumptions with biased outputs

All LLMs inherently produce hallucinations and biases based on factors such as misinformation, training limitations, and contradictory datasets. So, while genAI tools can play a valuable role in enhancing business decisions, their outputs must be augmented with human intelligence and the ability to check for accuracy, fairness, and overall fitness of purpose.

Suppose a startup decides to use a genAI tool to create messaging rather than hire a couple marketing copywriters. At face value, this might translate to sizable budget savings for a cash-strapped marketing leader. But it comes with a heavy downside as well: Without someone carefully reviewing and editing messaging before it goes out the door, the company runs the risk of the model misinterpreting the source material and expressing copy in a misconstrued or inappropriate way.

In some cases, the manpower required to check and fix these outputs ends up being just as time-consuming, if not more, than simply doing the work itself. If so, some companies may forgo the efforts altogether, accepting the risk (and potential consequences) of less-than-ideal brand messaging and quality.

#03

Four measures for keeping company data safe and secure

While the risks of data leakage cannot be overstated, it certainly doesn't mean that those using genAI applications are destined for data leaks or other problems. Often, it's not only what data you share but how you share it (and with whom) that keeps your data protected and secure.

To set your organization up for success, consider the following guardrails and best practices:

1. **Know what, how, and where your internal models are created via traceability and data lineage.**

   The first rule of data governance is you can't control what you can't see. This makes key capabilities such as data visibility, provenance, lineage, and traceability crucial to deploying any genAI use cases and applications.

   Managing hallucinations and biases becomes faster, easier, and far less time consuming as well. For example, inside the Vendia platform is an auditable, tamper-proof ledger that records data transactions across all systems, clouds, and parties. This makes it easy for business leaders to quickly trace outputs back to model sources and verify accuracy, as well as identify opportunities to regenerate models with updated or revised information.

2. **Keep a pulse on what data leaves (and stays!) inside your organization with full data visibility and fine-grained access controls.**

   Access to data lineage via an immutable transaction record also helps business leaders easily keep track of operational and analytical data exchanged between different departments or shared outside company walls—whether for genAI model training or other data-sharing purposes.

   This lets you see exactly where and how company data is being used across all operational, analytical, and external or third-party systems. Likewise, Vendia gives business leaders full control over who has access to their data at all times. Companies can choose what information to withhold or disclose based on individual roles, partners, records, and fields directly in the platform.

3. Enforce external guidelines around permissible data usage.

Data sharing between multiple organizations to create genAI models is becoming increasingly more common. But as powerful as these data-sharing networks may be, any data usage for LLM and genAI model training purposes will require careful, programmatic enforcement between parties.

Streamlining these parameters inside your data automation platform helps uphold data governance policies while reducing the burden of human auditors. Smart contracts inside Vendia can trigger workflows between systems or parties whenever business data is used in a defined way. Triggered at the point of transaction, enterprises gain the added benefit of enforcing these terms in real time, rather than waiting for an end-of-year audit…or an embarrassing headline.

4. Balance the need to share with the need to know.

Finally, approach all genAI deployments with care and caution, especially those involving PII and PHI. All companies are increasingly subject to tighter controls and regulations regarding consumer privacy, with compliance laws such as HIPAA, GDPR, and CCPA making it more challenging to store, copy, or share PII and PHI without the owner's agreement and consent.

While these regulations certainly aren't new, the ease with which folks can query data via genAI increases the risks of making compliance missteps without realizing it. Here, the easiest—and safest—path forward is using a data automation platform like Vendia that provides capabilities such as data redaction, masking, and anonymization, as well as automatic compliance and role-based access controls.

There's no getting around it: AI will only keep getting better and more adept at mimicking human interactions and intelligence. Which means that now is the best time to get the appropriate data sharing, handling, and governance safeguards in place to leverage a continuous wave of new innovation.

About Vendia

Vendia is the future of collective data intelligence, combining smart APIs, databases, and distributed ledger technology inside a single platform. Vendia's data automation cloud makes it easy to share data inside and outside of the organization in real time and with full visibility, governance, and control. Companies such as BMW, Delta Airlines, Resolution Life Insurance, and Fannie Mae use Vendia to automate contextual and compliant data flows between any-to-any systems for a harmonized, accurate view of data that unlocks speed, innovation, and cost savings. Learn more about us at Vendia.com and #UnchainYourData with Vendia.