



Using GenAI forecasting internally? Here's what to consider before scaling

02 April 2026, 14:15

Nick Boucart

Mihail Mihaylov

What if you go further with GenAI Forecasting?

Internal use of GenAI forecasting is both safe and valuable. However, once your application is exposed to the outside world, additional considerations come into play.

You can now build reliable forecasts using your own business data, without requiring ML expertise or expensive tools. But an important question remains: how safe is this approach?

For internal use: go ahead and apply

Are you using forecasting for internal analyses, management reporting or decision-making? Then the risks remain limited. Any errors stay within your organisation and are typically caught by people who understand the context.

The same applies to the models themselves. GenAI models generate predictions based on statistical patterns in data, because they do not truly understand what they generate. As a result, subtle errors can occur and may be difficult to detect for non-technical users. In an internal setting, these are more likely to be identified. Once the output is used externally, that becomes significantly more challenging.

Extra care is needed once

Extra care is required once the output of your model directly impacts third parties or feeds into automated processes. It is important to understand that GenAI models generate predictions based on statistical patterns, not on any real understanding of your processes or sector. As a result, incorrect predictions can still appear entirely plausible.

This is why a proper risk analysis is essential: what is the potential impact if the model is wrong, and how will you validate its output?

Consider, for example

- Predictions that automatically trigger orders or payments
- Forecasting output shared with customers or partners as factual information
- Applications where errors can lead to financial, legal or safety consequences

In these situations, a working application alone is not sufficient. You also need a structured approach to risk analysis, thorough testing procedures and a clear review process for model outputs.

What does this mean in practice?

The good news is that most time series forecasting applications start internally. Think of predictive maintenance for your own machines, monitoring energy consumption or optimising production planning. For these use cases, you can start applying what you have learned today.

As you move towards applications that extend beyond your organisation or become part of operational processes, new questions arise.

What are the risks if the model makes a mistake? Who is responsible for reviewing the output? And which safeguards need to be in place?

Sirris supports your next step

At Sirris, we work not only on implementing GenAI, but also on its reliability and robustness. We support companies with risk analysis, setting up testing processes and putting the right safeguards in place.

Want to know what the next step is for your application?

Contact our experts: [Nick Boucart](#) & [Mihail Mihaylov](#).

Would you like to help shape the future of digital product development with Generative AI?

Join the [LISA project](#) and collaborate with experts on the safe, robust and efficient integration of LLMs into your software products.

[Get in touch with project coordinator Nick Boucart by email](#)

Authors



Nick Boucart

Profile picture for user mihail.mihaylov@sirris..

Mihail Mihaylov