Arivonix AI
Data IN, Agentic AI ON, Revenue OUT

HOW A LEADING INSURER EMPOWERED UNDERWRITERS WITH AGENTIC AI WITHOUT REBUILDING THEIR SYSTEMS

> "Our underwriters are spending more time gathering data than assessing risk."

That's how the conversation began during the insurer's quarterly review. The underwriting head was visibly frustrated. Despite deploying decision support tools and digitizing parts of the workflow, turnaround times for complex policies had increased by 18% over the last two quarters. Worse, the most skilled underwriters were stuck performing low leverage tasks digging through lab reports, cross referencing old case files, and manually compiling risk notes.

An internal audit revealed the core issue: underwriters were spending over 60% of their time on data preparation alone. Manual errors had also crept in error rates in high volume policies had risen by nearly 15%, exposing the company to underwriting leakage and downstream compliance risk.

It was manual, slow, and frankly beneath their level of expertise.

Then came the real question from the COO:
“Why can’t we have AI agents that can do the heavy lifting for them?”

The vision was clear: underwriters should have intelligent agents that not only summarize data, but initiate actions to extract key info, flag inconsistencies, and suggest next steps. These agents wouldn’t just assist; they’d act. But that kind of agentic AI required a level of data readiness, model deployment, and regulatory traceability that the insurer simply didn’t have.

The CIO put it bluntly:
“Even if we had the smartest AI in the world, we don’t have a compliant, unified runway to let it operate.”

That’s where Arivonix stepped in.

The Use Case: Agentic AI for First Pass Risk Profiling

Rather than aiming for full automation, Arivonix AI proposed a high impact, low risk use case: let AI perform early-stage risk profiling for new applications before the human underwriter even begins reviewing.
The goal was simple: cut down time spent on routine data synthesis by having an agent pull together a “pre-underwritten view” using data from across systems medical records, past claims, policy history, and third-party lifestyle data.

But here’s the twist: the insurer didn’t want to move data around. Compliance was a top concern. The AI had to operate across sources without physically duplicating or transferring any sensitive data.

Using Arivonix’s bidirectional data mesh and virtualization, they connected 250+ sources of cloud warehouses, document repositories, internal databases without creating a new central lake. Data stayed in its original place, and the AI accessed only what it needed, when it needed it.

Training and Deploying the Agent

Inside Arivonix's DataForge Designer, the team trained an AI agent on a combination of historical underwriting decisions, domain rules, and real time ingesting from relevant sources. The agent learned to:

- Extract key values from PDFs and handwritten notes (like BMI, family history, blood pressure)
- Highlight missing documents or data points (e.g., medical exam not uploaded)
- Compare the application to similar historical profiles
- Generate a summary risk view with flags and confidence scores
- Most importantly, the model wasn't a black box. Every decision it made could be traced, explained, and controlled.

Guardrails, Governance, and Adoption

AI underwriting isn't just about performance, it's about trust. Arivonix built in the safeguards that allowed the insurer's legal and compliance teams to greenlight the use of agentic workflows:

- Role and permission-based access to AI results
- Subscription level visibility depending on the reviewer's clearance
- Audit logs that tracked every action the agent took and why

And the underwriters? They didn't lose control. Instead, they got a powerful shortcut:

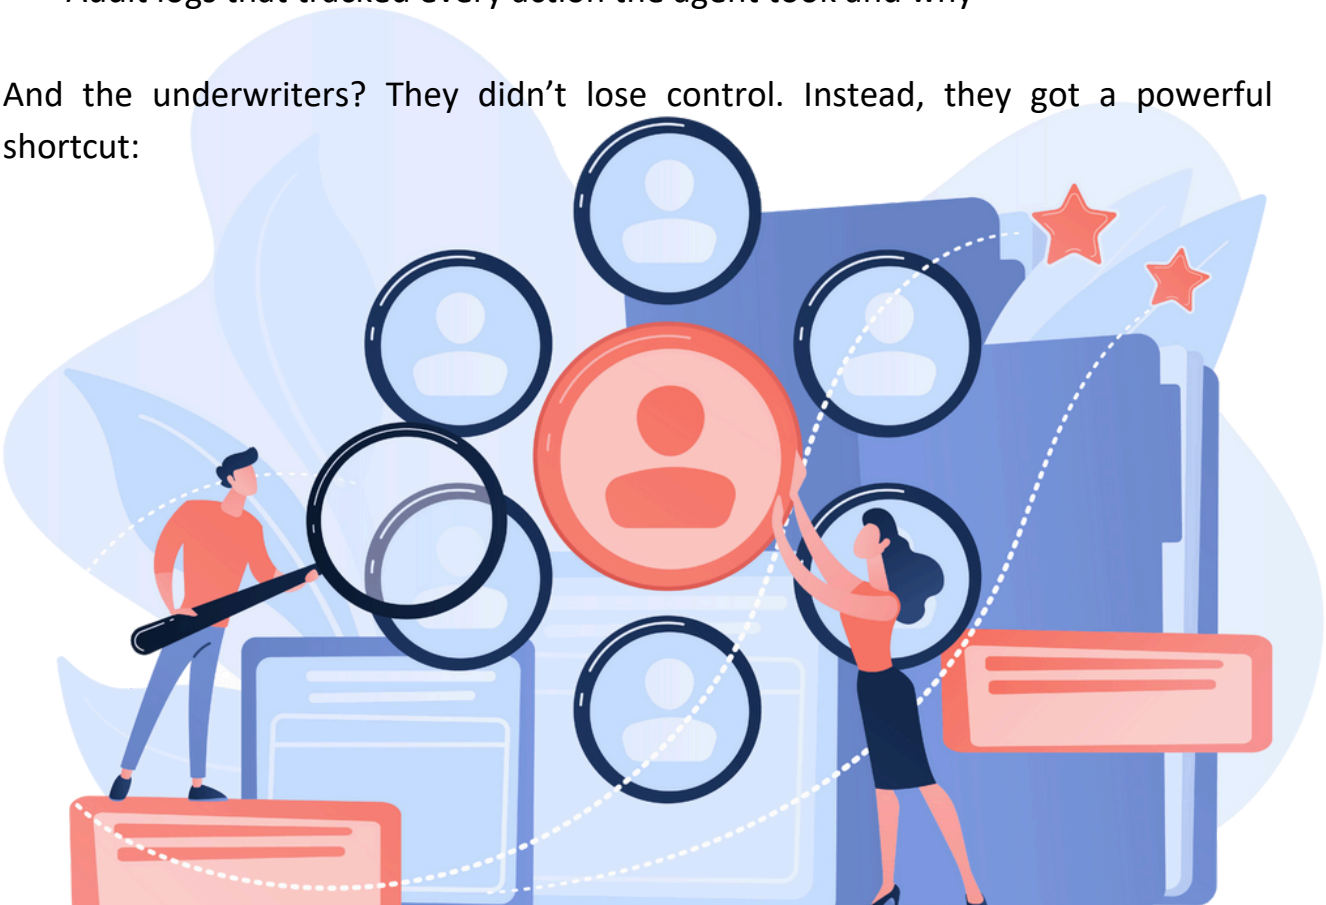

- Accept the AI generated profile
- Override it with notes
- Flag issues to retrain the model over time

It was still *their* decision. The agent simply tied it up.

The Impact

In less than 8 weeks:

- Underwriting prep time **dropped by 65%**
- Risk flag detection **improved by 21%**
- Error rates in high volume policies **fell by 40%**
- Most importantly, **adoption was near 100%** within the underwriting team

This wasn't a moonshot transformation. It was a pragmatic, secure, high value use case executed in a way that required no overhaul of legacy systems.

"We thought we needed to reinvent our tech stack to enable agentic AI. Arivonix showed us we just needed a smarter data fabric."
- Chief Underwriting Officer

Today, the insurer is expanding Arivonix agents to support claims triage, fraud pattern detection, and product personalization. But it all started with one use case: enabling underwriters to make faster, smarter decisions without ever losing control.

Conclusion:

A Blueprint for Scalable, Responsible AI in Underwriting
This case demonstrates that deploying Agentic AI in underwriting is not a future ambition; it's an achievable reality with the right data architecture. By anchoring the transformation around a specific, high value use case and leveraging Arivonix's unified data fabric, the insurer accelerated underwriting decisions, improved risk accuracy, and increased operational efficiency without disrupting existing systems.

Success underscores a critical lesson: intelligent automation doesn't require system overhauls. It requires a compliant, interoperable data foundation that enables AI to operate securely, transparently, and effectively within the business frameworks that already exist.