

No black boxes allowed: Secure, explainable AI for financial institutions

Ahead of the Curve, a Bankers Podcast

This transcript has been edited for clarity and readability.

Kate Randazzo:

This is Ahead of the Curve, a Bankers Podcast. On this episode, we're joined by Danny Piangerelli, vice president of engineering at Abrigo, where he leads data and AI platform engineering. With more than 30 years of experience in software and technology leadership, as well as previous roles as CTO at BankFi and Sensibil and co-founder and CTO of Malazai Software, Danny brings a rare blend of technical expertise and real-world fintech perspective.

In our conversation, Danny breaks down what responsible AI looks like in a regulated banking environment, why explainability and auditability are essential, and how Abrigo is building AI tools that empower people rather than replace them. From agentic frameworks such as AskAbrigo to the "Iron Man suit" analogy for innovation, this episode explores how financial institutions can adopt AI thoughtfully, securely and at a pace they can trust.

We've spent the last two years thinking about what responsible AI means for our organizations and how to stay safe if we decide to adopt it. Could you speak for Abrigo about how we ensure compliance and what explainability means to us?

Danny Piangerelli:

Those are excellent questions. Fundamentally, we approach any new technology choice, whether it's AI, agents or any other choice we've made over more than 20 years, from the position that Abrigo and our clients are regulated and audited. Nothing can be a black box.

In the world of AI, technology has traditionally been viewed as a bit of a black box. Questions get asked, magic happens, and something comes out the other side. The first things our customers need are data privacy and an understanding of where the "magic box" lives. They also need to understand what it is actually doing.

Our approach is the same one we use in most cases: We protect the data, protect the way we compute and access that data, and use the same formats we use everywhere else. Then we add explainability, giving systems the ability to explain and help audit what they are doing and the decisions they are making.

Kate Randazzo:

When we incorporate AI into decisions people are making within Abrigo solutions, how do we make sure they can see the back end? How is it auditable and understandable? Do they need to ask us for that information, or can they access it on their own within the solutions?

Danny Piangerelli:

The solutions are primarily built with auditability and explainability in mind. For example, in the fraud detection models we build for our AFD product, we have built the ability to explain how the models arrive at their answers directly into the models themselves. These are traditional machine learning predictive models.

The models actually produce an explanation. They can show, for example, "I'm weighting this particular value higher than this one, and therefore I'm arriving at this." We make that explanation available to the user in the application. Users can see the model and understand how it arrived at its result.

In other cases, when we are generating text through features such as narrative assistance, we provide documentation explaining where the data is coming from. In all of these cases, the models are not making a decision outside of a human-in-the-loop process.

In a narrative use case, for example, the system generates draft narratives. The user can review the draft, agree or disagree with it, edit it, delete it or do whatever they want with it. In our latest offering, AskAbrigo, we show when we pull documentation, data or explanations to answer questions. We include a reference to where the model got the answer. The user can click on each reference to see where the answer came from.

Kate Randazzo:

That's excellent. As someone who uses ChatGPT, my main use these days is taking Abrigo webinar transcripts and trying to build a GPT that can write in our tone and follow the Abrigo style guide. That way, I don't have to read a three-hour webinar transcript from beginning to end.

Unless I explicitly ask it to use only direct quotes from a webinar, it may summarize or infer beyond what the speaker said. So it's great to hear that our products can include a reference within each answer instead of forcing users to track down inconsistencies, read the source as closely and make sure that nothing has been hallucinated or assumed.

Could you go into a little more detail about AskAbrigo and how that agentic framework works? How is it different from tools such as ChatGPT?

Danny Piangerelli:

AskAbrigo is our knowledge-based agent. It started as a purely knowledge-based agent. When we met with customers, they told us they would love to access their customer data and use it in a ChatGPT-like way for things such as finding policy information or looking up how to train bank and credit union users on the ways we do things.

However, a bank or credit union cannot upload private documentation to ChatGPT. We gave customers a secure, Abrigo-hosted environment where their other applications and

data are hosted. They can upload documents safely and interact with those documents through the knowledge agent.

We call it an agent because we then added capabilities to it. In certain instances, the user can turn on internet activity. They can ask questions and conduct research on the internet, much like they might with ChatGPT, when the information is public and there is no concern about sharing private data.

Their data is still kept private and is not sent to the internet in any of these searches. We prevent the agent from accessing their internally uploaded data when they are interacting with the internet.

We've also rolled out new knowledge sources, such as a customer's own Sageworks data. If they are a Sageworks customer, they can ask questions about loans and borrowers and build views of that data.

All of that is what AskAbrigo does. It is built on top of a core agentic framework that we are now using to build other agents to do more within different areas of the business.

Kate Randazzo:

I've been impressed with how quickly we've been able to get these capabilities out. How do our internal practices enable us to scale this kind of development and expand our use of AI across our products effectively?

Danny Piangerelli:

We've certainly been on a journey with what we call our SDLC, or software development lifecycle. For many years, we had a traditional approach. AI tools came in and forced us to rethink a lot about how we build software.

It started a few years ago with using AI tools as an assistant alongside the developer. As people were writing code, they could ask the tool to double-check the code, look for bugs or do something similar.

It quickly moved into an agentic mode within these tools. We could ask the tool to build a set of tests, and then the developer could review those tests, decide whether they looked good, make modifications and send them along.

Since then, we've expanded this into the broader product lifecycle. Product managers are using tools to generate code even though they were not previously coders. They can generate user interface prototypes, product documentation and other materials. We are automating more of the product and ideation side of the process.

Now we're moving toward a process where we can take product ideas and interpret them through a flow that generates a product on the other side. Architects and developers can modify, augment or delete the generated code, but they are not starting from scratch.

The developers and architects still own and review the code. It remains a focused project, but they don't have to begin with a blank screen. We've seen a tremendous increase in the amount of work we can get done as a result of these changes.

Kate Randazzo:

That makes perfect sense. I have a writing degree and no finance background. When I came to Abrigo, the company could have spent a lot of time and resources teaching me every aspect of banking, or it could give me someone like Dave Koch and ask me to turn what he said into a blog.

That sounds similar to what AI can do for coders, developers and product teams. Instead of starting from scratch, someone without a coding background can describe what they want an idea to look like and then fine-tune the result. Is that a fair metaphor for what's happening for people who are unfamiliar with AI?

Danny Piangerelli:

Yes, that's exactly right. In software development, one area where we're seeing a significant change is documentation.

Traditionally, an architect would draw diagrams to show how the software works. Within a week or two, those diagrams could already be out of date because we had made design changes or decided to change the way something worked. The architect often was not going back to update the documentation.

Now we are generating that documentation with these tools, and the tools generate it on the fly. If someone makes a code change that affects the architecture of the project, the tools can generate a new architecture diagram in real time. The diagram shows what the application is actually supposed to look like.

That saves a significant amount of time. It also produces documentation that is up to date and accurate because we are not relying on someone to go back and make manual changes.

Kate Randazzo:

That's remarkable. It's great that we can make this kind of design work more accessible to people at Abrigo who have good ideas or hear good ideas from our customers.

Are there other ways we're investing in team development and supporting this kind of innovation, especially since the technology is moving so quickly?

Danny Piangerelli:

There is a lot of discussion in the industry about what happens to roles that have traditionally been responsible for this work as we generate code and use tools to perform some of these tasks.

Abrigo has not primarily taken a replacement approach. We are not trying to replace developers, product teams or marketing teams with tools that generate content or code. Instead, we have taken what we call the Iron Man approach.

Instead of building a robot that goes and does your job, what if we build an Iron Man suit and put that suit on you? You are still in control, but you are empowered to do things you could not do in the past.

What we're seeing now is that some development roles are becoming more focused on higher-level thinking and leadership. A developer can still focus on writing specific code, but that person is also empowered to think at the level of an architect or a team lead, orchestrating a set of tools to do work.

We're trying to teach people to think of themselves as the conductor of an orchestra of agents that can go off and build things for them.

As you said, building a GPT forces you to think differently. Instead of asking only how to get a tool to do something, you have to think about the output you actually need. You also have to communicate well. If you leave out an important detail or do not explain something clearly, the tool can misunderstand what you mean.

We are investing in helping our teams communicate better and think at a higher level, so they can operate more like orchestrators or architects.

Kate Randazzo:

I like that comparison to conducting, and we use it often when we explain the value of these products to customers. When relationship building is your bread and butter, especially in smaller communities, freeing up time to build relationships can be incredibly valuable.

It also gives people a chance to think about their workflows in a way they may not have had to before, especially if they've done things the same way for the last 10 or 20 years. They can think about what they want their workflows to look like, what they want their free time to look like and what they want more time to do in person, including the things AI cannot replicate.

That exercise can be useful even if it does not lead immediately to a GPT or another AI tool. Simply thinking about what you want your output to be in your job can be valuable.

Danny Piangerelli:

When you're given a kind of supercharger, instead of viewing it as a threat, think about the difference between having a shovel and having a machine that can dig up all this dirt. Now you can do much more with these tools. Let them expand what you think you are capable of doing.

Kate Randazzo:

Very Iron Man.

Switching gears a little, what progress has Abrigo made in enabling customers to self-serve or purchase AI tools online themselves?

Danny Piangerelli:

We are moving in that direction. AskAbrigo, for example, started as a standalone product. It did not require an implementation. Customers could sign up for the functionality, and during the beta period we gave them free access. They could sign up and get access without any additional implementation work.

As we think through other features we've added, such as narrative assistance, these are features that are added to the normal application. When they become available, they appear in the latest version and are turned on.

For every decision we make about enabling these capabilities in our products, we are trying to determine the path of least resistance. Even if we reimaged an entire loan origination system, we would need to consider how to avoid a massive migration or a difficult old-versus-new discussion.

That is part of our design thinking around enabling these capabilities: We want to create the least amount of friction possible. We are still learning.

Kate Randazzo:

That's wonderful. I have one last question for you, and it might be a little broad. What advice would you give to other fintech engineering teams that are trying to build AI programs in regulated environments like ours?

Danny Piangerelli:

There is a lot changing, and there is a lot to consider. I think people in our industry can take one of two extreme paths. We are trying to go down the middle.

One extreme is to ignore AI and say it is hype or the latest trend, and that we should rely only on tried-and-true practices. I think that is the wrong approach because the train has effectively already left the station. People are not debating whether AI will change things. They are debating how much it will change things.

There is no question that AI will change the way we build software and run our businesses. Ignoring it completely is the wrong approach.

The other extreme is assuming that every part of the business and every way of doing work will somehow need an AI component. Becoming hyperfocused on learning the new thing is also an overreaction.

Our customers need products that work and continue to work. If we take our eye off that priority and become a research shop that builds experimental products simply because they are the latest thing, that is another overreaction.

What we are trying to do is start with a foundation and build from there. We have built a business on top of a secure, resilient underlying data structure that meets regulations and passes audits. On top of that, we have built software that our customers are pleased with. It is growing, getting better and getting stronger.

The question is how we build on those foundational elements and decide which initial bets to make with AI without losing sight of them. As we said at the beginning, observability, explainability and an understanding of what the tools are doing matter regardless of whether we are building AI-enabled products.

Customers need transparency and need to understand what a product is doing. We are focused on the latest and greatest technology and on incorporating it at a pace our customers can depend on and trust while still staying out in front.

That balance is hard to achieve. It is hard to balance.

Kate Randazzo:

I'm glad we have you at the helm as we try to find that balance. I appreciate your time and expertise, and we're looking forward to hearing what Abrigo comes up with next in the AI space. Hopefully we'll have you on again.

Thank you.

Danny Piangerelli:

Thank you.

Kate Randazzo:

You can find this and future episodes of Ahead of the Curve on [Abrigo.com](https://abrigo.com) or on your favorite podcast app or platform. Search for Ahead of the Curve, a Bankers Podcast, or search for Abrigo, and subscribe so you don't miss future episodes.

Thank you so much for listening. We'll be back with our next episode as soon as we can.