



AI in Collateral Management Where It Adds Value and What It Takes to Make It Work

More Expertise. **More** Trust. **More** Achieved.

Where AI adds value, what it needs to work and how firms can deploy it safely.

Contact Us



The Tonic Consultancy
info@thetonicconsultancy.com

CONTENTS

01	Three types of AI: why the difference matters.....	3
02	Applying AI to the collateral process	4
03	When AI adds value and when rules are better	5
04	Connecting AI to collateral systems.....	6
05	Control should tighten as AI moves towards action.....	7
06	What this means for collateral	8
07	How Tonic can help	9

AI in Collateral Management - Where It Adds Value and What It Takes to Make It Work

Philip Forkan | Head of Clearing and Collateral, The Tonic Consultancy

Reading time: ~7 minutes

Every few years, someone says collateral management is about to change. Over twenty-five years in margin and collateral, I have heard the same claim around post-crisis clearing, the uncleared margin rules under EMIR and Dodd-Frank, and each new wave of automation. Most changes made the same function faster: fewer manual steps, better controls and quicker processing.

AI adds something different to conventional automation. It can interpret unstructured information, forecast changing conditions and support decisions where the answer is not fixed. Its value, however, depends less on the model than on the data, system access and controls around it.

The World Economic Forum's June 2026 AI Playbook for Financial Services draws on more than 150 senior leaders across more than 100 organisations. As financial services moves beyond experimentation, it emphasises governance, data and infrastructure. For collateral teams, those foundations often determine whether an AI use case can move from pilot to production.¹

That makes the distinction between predictive, generative and agentic AI important: they solve different problems and require different levels of control.

01

Three types of AI: why the difference matters

AI is often used as a single label, but predictive AI, generative AI and agentic AI do different jobs. The distinction matters because they use different inputs, operate at different levels of maturity and require different controls.

Table 1: Three types of AI, with different levels of maturity

	PREDICTIVE AI			GENERATIVE AI			AGENTIC AI		
Maturity today	In production	Piloting	Frontier	In production	Piloting	Frontier	In production	Piloting	Frontier
What it does	Identifies patterns, predicts outcomes, optimises decisions			Interprets and creates content from unstructured data			Uses tools to carry out approved actions within controls		
Collateral application	Liquidity forecasting, inventory analytics, exposure prediction			Legal agreement processing, dispute analysis, counterparty correspondence			Margin-call and substitution workflows, exception handling and controlled execution		

Predictive AI has been used for longer in areas such as forecasting, pattern detection and decision support. Generative AI interprets and creates content, particularly from unstructured information such as documents, emails and correspondence. Agentic AI goes further by using tools and taking actions within defined limits. For many practitioners, the most visible recent change has been generative AI through desktop tools used to summarise documents, draft text or answer questions, but this is only one part of the AI landscape.

02


Applying AI to the collateral process

Across the collateral lifecycle, predictive AI can anticipate liquidity or collateral needs; generative AI can interpret agreements and margin correspondence; and agentic AI can use approved tools to carry out defined actions.

These capabilities should complement the rules, workflow tools and optimisation engines already used in collateral management. A conventional optimiser, for example, can identify the cheapest eligible asset when the inputs and constraints are known. AI becomes more useful when the decision also depends on uncertain or unstructured information, such as forecast liquidity needs, changing agreement terms or market conditions.

Figure 1 maps those capabilities across the collateral process and the systems firms use today. It shows where AI could apply.

Figure 1: Applying the three AI types across the collateral process



03

When AI adds value and when rules are better

Much of collateral management is deliberately rules-based. Margin requirements are calculated against agreed methodologies, collateral must satisfy CSA or CCP eligibility rules, concentration limits and haircuts, and settlement instructions have to meet defined cut-offs. Where the inputs and constraints are known, conventional workflow tools and optimisation engines can usually produce the right answer more reliably and transparently than AI.


The harder decisions sit around those processes. A collateral manager may need to choose which eligible asset to post while preserving cash for other obligations, or whether to substitute securities that have become expensive to finance. They may also need to decide how much inventory to retain against expected margin calls or how to respond when settlement capacity or market conditions change during the day. These decisions combine hard constraints with forecasts, market information, funding costs and operational judgement.

That is where AI can add something different. Predictive models can help anticipate margin and liquidity needs. Generative AI can interpret agreement terms, margin correspondence and other unstructured information. AI-supported decision tools can bring those signals together when comparing collateral or liquidity options. The existing eligibility and settlement controls should remain deterministic.

The test is whether AI improves a decision that existing rules or optimisation cannot solve on their own. If a task has one defined answer, use the rule. If a conventional optimiser can identify the lowest-cost eligible asset from known inputs, use the optimiser. AI becomes more relevant when the decision depends on uncertain future needs, changing funding or market conditions, unstructured information, or several competing objectives.

Margin processing illustrates the boundary. Calculating and reconciling a standard margin call is largely deterministic. Deciding how best to meet it may require a wider view of eligible inventory, expected liquidity, funding costs and settlement cut-offs. The clearest generative-AI opportunity is often in the exceptions: interpreting correspondence, structuring information from emails and attachments, identifying the likely cause of a dispute and suggesting the next action while users retain control. Figure 2 applies the same test across the collateral process.

Figure 2: Testing which collateral tasks actually benefit from AI



04


Connecting AI to collateral systems

Collateral teams rarely make decisions from one system. Understanding how to meet a margin call may require inventory, agreement terms, eligibility rules, prices, funding costs, liquidity needs and settlement status, often held across different platforms and reports. AI can only support that decision if it can access and combine the same current information in a controlled way.

Model Context Protocol, or MCP, is one emerging open standard for doing this. In simple terms, it gives an AI application a common way to find and use approved data and tools. Existing collateral platforms and APIs still do the underlying work; MCP gives AI a consistent, controlled way to access the data and functions they already provide.ⁱⁱ

For a collateral user, the benefit is less swivel-chair work and a shorter path from question to action. A user could ask what is driving a margin requirement, which eligible assets are available, which positions are expensive to post or which movements are at risk of missing settlement. Where permissions allow, an agent could also assemble the information needed for a substitution or prepare a defined workflow action for validation. Figure 3 shows that flow from user question to explanation, recommendation or prepared action.

Figure 3: How AI can work across existing collateral systems



Access still needs clear boundaries. The underlying collateral systems remain authoritative, and existing user permissions, eligibility rules and approval controls should determine what AI can see and do. Reading data, recommending an action, preparing a workflow item and executing it should be treated as separate levels of authority.

05

Control should tighten as AI moves towards action

In collateral, governance becomes real at the point where an AI output can change a position or move an asset. The key question is what happens when the model is wrong: an incorrect interpretation can lead to an ineligible asset, unnecessary use of cash, a missed settlement cut-off or the wrong movement being prepared.

That argues for a firm boundary between AI judgement and deterministic control. AI may interpret a CSA, investigate a dispute, forecast a margin requirement or recommend an asset. Before any recommendation becomes an action, existing controls should re-check eligibility, haircuts, concentration limits, credit and liquidity constraints, settlement cut-offs and approval thresholds.

Authority should increase in stages. Read-only access can be relatively broad; recommendations can go further where they are revalidated; preparing a movement or response needs tighter checks; and direct execution should be limited to defined actions, systems, counterparties, values and time windows. Figure 4 sets out that progression from Read to Act.

Figure 4: Increasing control as AI moves from Read to Act



The most important controls are around failure modes the collateral desk already understands: stale inventory or prices, misread agreement terms, ineligible collateral, concentration breaches, liquidity already committed elsewhere, duplicate instructions and missed settlement cut-offs. These risks already exist, but AI can propagate a bad input or decision faster, so source data, validation rules and stop conditions matter.

High-value or unusual movements should retain existing maker-checker and approval thresholds. The audit trail should cover the decision as well as the final transaction: what the AI saw, what it recommended, which rules were applied, whether a person approved it and what was ultimately executed.

AI can broaden the information used in a collateral decision, but the closer it gets to moving assets, the narrower its authority should become.

06

What this means for collateral

AI should complement rules, optimisation engines and experienced collateral teams. Its value is highest where the answer depends on uncertain, changing or unstructured information and still has to stay within hard constraints.

The opportunity is to connect those capabilities to existing workflows: use AI to interpret, forecast and recommend; use established rules, permissions and approved tools to validate and execute.

The harder work sits underneath. Firms need current data, clear rules, controlled access to systems and a record of why an action was taken. Without those foundations, AI simply makes a weak process faster.

The immediate question for collateral teams is which decisions are worth improving, what data they require, and how much authority the firm is prepared to give AI.


07

How Tonic can help

Tonic helps firms turn those questions into a delivery plan. We assess data, systems, controls and operating processes, identify the use cases with a credible business case, and define the changes needed before AI can move from analysis into controlled production use.

The work can start small. The aim is to get value from useful applications now while fixing the data, system links and controls that more advanced use will require. Figure 5 sets out six steps.

Figure 5: Six practical steps



Philip Forkan leads Clearing and Collateral at The Tonic Consultancy. Tonic advises financial institutions on collateral, clearing and the practical use of AI.

ⁱ World Economic Forum, The AI Playbook for Financial Services, 24 June 2026, [weforum.org](https://www.weforum.org).

ⁱⁱ Linux Foundation, “Linux Foundation Announces the Formation of the Agentic AI Foundation”, 9 December 2025, [linuxfoundation.org](https://www.linuxfoundation.org).