C&R Software
Credit Risk Management

White Paper

Closing the AI execution gap in collections

Banks have spent years talking about AI transformation in collections. The investment is real, and the urgency is growing. But most institutions are still running proofs of concept rather than production ready capabilities. This gap between experimentation and scale isn't a technology problem. It's an architecture problem.

Point solutions can prove AI's value in isolation. What they can't do is give collections teams a governed, reusable foundation for deploying AI across the full lifecycle. That's what an agentic framework provides. It embeds AI inside the collections environment where the work already happens, connects data, decisioning, and compliance controls in one place, and makes every new use case faster to build, test, and deploy.

This paper explains why the point solution model creates barriers to scale, what an agentic framework does differently, and what responsible AI deployment in collections looks like when governance is built in from the start.

The gap between AI investment and impact

The promise of AI in collections is well-established. Better treatment strategies, more consistent customer interactions, smarter contact decisions, reduced manual effort, and stronger compliance outcomes are all within reach. Banks know this. That's why AI investment has accelerated in recent years.

But investment and impact aren't the same thing. Despite rising budgets and growing ambition, Deloitte's **2026 Banking and Capital Markets Outlook** found many AI initiatives remain stuck in isolated proofs of concept, held back by weak governance, duplicated effort, and uneven results across the organization. The consequence is predictable: when governance can't keep pace with deployment, projects stall. **46% of AI initiatives** are abandoned between proof of concept and broad adoption, not because the technology failed, but because the operating model wasn't ready for it.

Collections teams know this pattern well. Most remain in experimentation mode, testing tools in contained environments and cautiously expanding scope while waiting for clearer paths to production. The bottleneck isn't a shortage of use cases or a lack of willingness to invest. It's the absence of an architecture that can take AI from promising pilot to governed, scalable capability.

All content produced by C&R Software for marketing purposes is exclusive to the company's marketing campaign. Unauthorized use, reproduction, distribution, or transmission in any form without prior written permission from C&R Software is strictly prohibited.

© 2023–2026 Jonas Collections & Recovery Inc, ("C&R Software") - All Rights Reserved.
C&R Software is a subsidiary of Constellation Software, Inc.

Why point solutions create new problems

When a collections challenge surfaces, the instinct is to find a solution built specifically for it. One vendor for propensity scoring. Another for agent guidance. Another for outreach optimization. Another for treatment segmentation. Each one works in isolation. But together, they create a fragmented AI environment that's hard to govern, scale, and sustain over time.

Data fragmentation

Collections is inherently contextual. The right treatment for a customer depends on account status, payment history, communication preferences, vulnerability indicators, policy rules, and real time operational signals. When AI tools operate across separate data layers, each system sees only part of that picture. **Over 90% of data users at banks** say the data they need is often unavailable or slow to retrieve, and 81% cite data quality as a top challenge to AI adoption. Point solutions don't fix that. They entrench it.

Governance fragmentation

Every AI model deployed in a regulated financial institution carries documentation, validation, monitoring, and auditability obligations. Traditional model risk frameworks are increasingly seen as **insufficient for AI**, demanding a shift from point-in-time controls to continuous, risk-based governance. This challenge compounds when models are spread across multiple vendors and platforms, each with its own governance artefacts and oversight requirements.

Vendor sprawl

Every additional point solution brings another contract, another integration, another security review, another data processing agreement, and another operational dependency. Collections leaders are increasingly asking a practical question: are we building a coherent AI capability, or are we building a vendor management problem?

Limited reuse

Point solutions are designed to solve defined problems. They're not built to be combined, coordinated, or extended into adjacent use cases without significant re-integration effort. That means every new collections AI capability risks triggering another full procurement and deployment cycle, with no compounding return on the infrastructure investment already made.

All content produced by C&R Software for marketing purposes is exclusive to the company's marketing campaign. Unauthorized use, reproduction, distribution, or transmission in any form without prior written permission from C&R Software is strictly prohibited.

© 2023–2026 Jonas Collections & Recovery Inc, ("C&R Software") - All Rights Reserved.
C&R Software is a subsidiary of Constellation Software, Inc.

A different starting point: The agentic framework

An agentic framework takes a fundamentally different approach. Rather than assembling individual AI tools around specific problems, it provides a unified environment where AI agents can be built, tested, orchestrated, and deployed across any collections use case inside a single governed infrastructure.

This matters because collections workflows are multi-step and deeply contextual. Customer contact doesn't exist in isolation. It sits within an account history, a treatment strategy, a regulatory context, and a portfolio risk profile. Effective collections AI needs to reason across those dimensions simultaneously, not hand off between disconnected systems.

An agentic framework makes this possible. It creates a foundation where AI capabilities compound rather than fragment, where new use cases inherit the same data, compliance, and governance infrastructure already in place, and where the distance between pilot and production gets shorter with every deployment.

The AI framework creates shorter integration timelines, lower deployment risk, and a governance baseline that doesn't need to be rebuilt for every new use case.

How the C&R Agentic Framework works

C&R Software's Agentic Framework is designed to enable the creation, testing, and deployment of AI agents within C&R's own infrastructure. It connects pre-built AI agents and Debt Manager functions into custom AI agents that can be configured for any use case within the platform.

The framework is built on Amazon Bedrock, which supports multi-model, multimodal AI that brings together machine learning, large language models, and generative AI across the C&R technology stack. The architecture uses an orchestration layer to coordinate specialized agents handling data retrieval, filtering, reasoning, and decision support, each drawing on knowledge bases that include internal documents, external data, and user feedback.

Critically, the framework is deployed within C&R Software's cloud infrastructure, which is already approved and in use at major financial institutions running Debt Manager in the AWS US-East-1 region. This means the AI operates where the collections data already lives, inside a trusted, regulated environment rather than sitting alongside it. The result is shorter integration timelines, lower deployment risk, and a governance baseline that doesn't need to be rebuilt for every new use case.

All content produced by C&R Software for marketing purposes is exclusive to the company's marketing campaign. Unauthorized use, reproduction, distribution, or transmission in any form without prior written permission from C&R Software is strictly prohibited.

© 2023–2026 Jonas Collections & Recovery Inc, ("C&R Software") - All Rights Reserved.
C&R Software is a subsidiary of Constellation Software, Inc.

Collections use cases the framework can support

Because the framework operates inside the Debt Manager environment, it can be configured across the collections lifecycle without requiring separate integrations or standalone deployments.
Use cases include:

Agent guidance and real time support.
During live customer calls or account work, AI can surface policy-compliant suggested responses, escalation prompts, and next-step guidance based on approved institutional documentation.

Treatment strategy and segmentation.
Specialized agents can retrieve account data, assess propensity signals, apply policy logic, and recommend treatment pathways, all within a single coordinated workflow rather than across disconnected tools.

Portfolio monitoring and exception detection.
Agents can run continuous reviews against defined thresholds, flag accounts requiring human review, and escalate exceptions through the appropriate workflow steps.

Post-interaction documentation.
AI can support automated logging, call summarization, and outcome recording, reducing manual workload and improving consistency of records for audit and compliance purposes.

Quality assurance and compliance review.
Logged AI suggestions and agent actions create a complete interaction record that supports both internal QA processes and regulatory examination requirements.

Each of these capabilities operates inside the same governed environment, drawing on the same knowledge bases and inheriting the same audit, privacy, and oversight controls.

Governance built in, not bolted on

For regulated financial institutions, governance is often the difference between an AI pilot that looks promising and an AI capability that's actually approved for production. The more AI expands, the more that distinction matters.

Institutions need to move from periodic, checklist-driven controls to continuous, risk-based oversight that covers development, validation, monitoring, and third-party AI. The Wolters Kluwer **Q1 2026 Banking Compliance AI Trend Report** found that explainability and transparency were cited as the most acute regulatory concerns by financial institutions, and that banks advancing AI without concurrent governance frameworks risk regulatory scrutiny from model risk management examinations and fair lending reviews.

All content produced by C&R Software for marketing purposes is exclusive to the company's marketing campaign. Unauthorized use, reproduction, distribution, or transmission in any form without prior written permission from C&R Software is strictly prohibited.

© 2023–2026 Jonas Collections & Recovery Inc, ("C&R Software") - All Rights Reserved.
C&R Software is a subsidiary of Constellation Software, Inc.

The C&R Agentic Framework addresses this directly. Governance isn't applied after deployment. It's structural.

- Human-in-the-loop design. AI presents recommendations. Human agents retain full responsibility for final decisions and customer interactions. There are no autonomous customer-facing actions.
- Policy-driven outputs. All AI-generated suggestions are grounded in institution-approved documentation. The system doesn't draw on the public internet or uncontrolled external sources.
- Audit trails by default. Every interaction is logged, creating a complete and traceable record of AI suggestions and agent actions for compliance review and regulatory examination.
- Privacy by design. The framework doesn't access, process, or store personally identifiable information or customer financial data beyond its designated scope. Strict controls within AWS Bedrock ensure knowledge bases and user interactions aren't used for model training purposes.
- Data segregation. Controls within the AWS Bedrock component prevent knowledge bases and user interactions from being used to train the underlying AI agents.
- LLM flexibility. The default model is selected through testing as the most effective for each use case, but clients can choose a different underlying model if their governance requirements or preferences call for it.

Together, these controls make it possible to scale AI use cases without scaling governance complexity at the same rate. As new agents are added within the framework, they inherit the same oversight architecture rather than requiring a separate governance build.

The strategic shift

The collections industry doesn't have an AI idea problem. It has an execution problem. Growing concern about vendor proliferation, governance gaps, and fragmented data environments reflects a deeper question that many institutions are now asking: are we building a coherent AI capability, or are we accumulating a set of tools that are individually useful but collectively unmanageable?

The institutions that close the AI execution gap won't necessarily be the ones running the most pilots. They'll be the ones that build the operating model required to take AI from experimentation to production reliably, repeatedly, and at scale. The answer isn't speed at the expense of oversight, and it isn't governance at the expense of progress. It's an architecture that enables both.

An agentic framework provides exactly that. It replaces a tool-by-tool approach with a governed environment for building and deploying AI where collections work actually happens. Every new use case is faster to deploy, every governance obligation is easier to meet, and every investment in the framework compounds across future capabilities rather than standing alone.

That's what responsible AI at scale looks like in collections.

C&R Software has delivered collections and recovery technology to financial institutions globally for over four decades. The C&R Agentic Framework is built within the same trusted infrastructure that powers Debt Manager deployments at major banks worldwide.

To learn more, visit **crsoftware.com**

All content produced by C&R Software for marketing purposes is exclusive to the company's marketing campaign. Unauthorized use, reproduction, distribution, or transmission in any form without prior written permission from C&R Software is strictly prohibited.

© 2023–2026 Jonas Collections & Recovery Inc, ("C&R Software") – All Rights Reserved.
C&R Software is a subsidiary of Constellation Software, Inc.