Rocket® software

Agentic AI in Core Banking Operations

Why agentic AI can increase risk in banking

Introduction

Agentic artificial intelligence is rapidly moving from an advisory role into direct execution. For banks this transition represents a fundamental shift in how IT operations are managed. Banking mission-critical systems are tightly coupled, heavily regulated, and completely intolerant of error. When you introduce autonomous agents that can execute actions within these environments, the stakes change entirely.

The cost of failure in a core banking system isn't just an IT inconvenience. A misstep, such as an AI agent misinterpreting a transaction command, can lead to immediate customer impact, severe regulatory exposure, and lasting reputational damage. While AI promises incredible efficiency and speed, deploying it without strict controls introduces unacceptable risk. In banking, automation does not fail quietly. It fails systemically.

Enterprise leaders must recognize that adopting agentic AI is no longer a conversation about the intelligence of the model. It is a conversation about operational trust. To modernize without disruption, banks must prioritize governance over raw capability.

Where agentic AI breaks in core banking environments

To understand the risks of ungoverned AI, we must look at how core banking systems actually operate under pressure. These environments process millions of transactions daily, relying on highly orchestrated workflows that leave no room for guesswork.

When agentic AI acts without a complete understanding of the system context, it amplifies the risk of cascading failures. Consider these critical operational moments:

01

Nightly batch execution and ABEND storms

Batch processing is the heartbeat of a bank. If an AI agent attempts to automatically remediate an abnormal end (ABEND) without understanding the downstream dependencies, it can corrupt data ledgers and delay market openings.

02

End-of-day settlement and clearing windows

These critical windows operate on strict service-level agreements. An ungoverned automated action that slows down processing can cause a bank to miss clearing deadlines, resulting in massive financial penalties.

03

Peak payment processing

During high-volume periods for ACH, wire transfers, and real-time payments, the network must perform flawlessly. An agentic tool that misconfigures a network setting in an attempt to optimize traffic can inadvertently halt transactions.

04

Cross-system incidents

Modern banking applications span distributed cloud environments and legacy mainframes, touching core applications, Db2® databases, CICS transaction servers, and security layers. An AI acting in a silo cannot safely resolve incidents that span these complex architectures.

> Agentic action without full system context does not solve problems. It creates them.

The banking risk multiplier: regulation, complexity, and the skills gap

The operational risks of agentic AI are multiplied by the structural pressures unique to the financial sector.

First, regulatory scrutiny demands complete transparency. Regulators require strict traceability, auditability, and explainability for any action taken within a core system. If an AI agent executes a change, you must be able to prove exactly why it happened and who authorized it.

Second, banking systems are incredibly complex. They have evolved over decades, resulting in hybrid environments that depend heavily on institutional knowledge. The logic behind certain configurations is often undocumented, living only in the minds of veteran engineers.

This brings us to the third multiplier: the widening skills gap. As deep subject matter experts retire, the fragility of these systems increases. The traditional manual war-room response, where dozens of experts gather to diagnose a cross-system incident, no longer scales. There are simply not enough experts left to manually untangle every complex outage.

Banks do not lack automation. They lack governed automation that can safely bridge this skills gap.

What a safe AI operating model looks like for banks

To safely introduce execution into core operations, banks need a robust AI operating model. This model must prioritize safety, compliance, and deterministic execution over speed.

These are not AI differentiators. These are fundamental banking requirements.

Policy-as-code aligned to bank controls: Agentic actions must be governed by strict policies that are hardcoded into the execution layer. The AI must automatically check these policies before taking any action, ensuring compliance at all times.

Deterministic execution: The system must validate every proposed action before it touches a production environment. There can be no hallucinations or unauthorized experimentation.

Identity-aware action authorization: Execution permissions must be role-based and tied to specific, verifiable identities. The AI should only act within the exact permissions granted to it.

Autonomous stopping capability: If the AI encounters a scenario where confidence is low or the context is unclear, it must have the ability to stop execution autonomously and escalate to a human expert.

End-to-end diagnostic trace: When diagnosing an issue, the AI must be able to trace the problem comprehensively from the initial symptom, through the transaction and job, down to the responsible root code.

Immutable audit trails: Every action, decision, and policy check must be logged securely in a tamper-resistant format that satisfies both internal compliance teams and external regulators.

Day-one value: how banks can adopt agentic AI safely

Adopting agentic AI does not require a risky, all-at-once deployment. Banks can achieve day-one value through a pragmatic, phased approach that builds trust over time.

1. **Read-only diagnostics:** Begin by deploying the AI to monitor systems and trace issues without the ability to make changes. This allows teams to validate the AI's diagnostic accuracy safely.
2. **Automated correlation and root cause analysis:** Allow the AI to automatically group alerts, filter out noise, and identify the root cause of complex incidents, drastically reducing the time teams spend investigating issues.
3. **Governed recommendations:** The AI proposes specific remediation actions based on its diagnostics, but a human operator must review and approve the action before execution.
4. **Selective execution under policy:** Once trust is established, authorize the AI to execute specific, low-risk actions automatically, strictly governed by policy-as-code.

This phased approach empowers banks to tackle high-value use cases immediately. For example, teams can automate batch failure triage, conduct rapid SLA impact analysis, and perform deep root-cause investigations for cross-system performance issues, all while maintaining absolute control.

Closing: why this is an operating discipline problem

The successful integration of artificial intelligence into banking infrastructure requires a fundamental shift in perspective. Banking success with agentic AI depends less on model intelligence and more on operating discipline.

You must ensure that any automated agent inherits the same rigorous controls, accountability, and compliance standards that you demand of your human operators. When you build guardrails directly into the fabric of your AI operating model, you transform a potential risk into a powerful competitive advantage.

Platforms designed for governed, operational-grade agentic AI, such as Rocket® EVA™, are built specifically to support this discipline in core banking environments. By focusing on deep diagnostics and governed execution, you can empower your team to operate securely, ensuring that your enterprise innovates confidently on the systems that power your business.

About Rocket Software

Rocket Software is a global technology leader in modernization and a partner of choice that empowers the world's leading businesses on their modernization journeys, spanning core systems to the cloud. Trusted by over 12,500 customers and 750 partners, and with more than 3,200 global employees, Rocket Software enables customers to maximize their data, applications, and infrastructure to deliver critical services that power our modern world. Rocket Software is a privately held U.S. corporation headquartered in the Boston area with centers of excellence strategically located throughout North America, Europe, Asia and Australia. Rocket Software is a portfolio company of Bain Capital Private Equity.

Rocket® software

Modernization. **Without Disruption.**™

Visit RocketSoftware.com

MAR-18235_WP_EVAMidFunnelExecBrief_V2