Bretton AI

The Compliance Leader's Guide to Agentic AI

bretton.com

The status quo for financial services compliance is needlessly costly—and exhausting. Large institutions spend an estimated $500 million to $1 billion annually on analysts who perform repetitive, manual compliance work, often clearing alerts that shouldn't have been flagged in the first place.

This substantial expenditure reflects a calculated risk mitigation strategy. By implementing comprehensive alerting protocols, the thinking goes, compliance teams are more likely to catch potential indicators of risk exposure.

But the overly cautious route creates unnecessary work. With so many alerts to investigate, analysts spend significant time on clearing cases that shouldn't have been flagged, leaving less capacity for complex investigations that require more effort.

MIT research validates the efficacy of AI for back-office automation, which offers some of the most dramatic and sustainable return on investment. But when it comes to finding the right agentic AI solution, it can be difficult to know what to look for. Solutions vary widely in their architecture, capabilities, and regulatory readiness. To help enterprise buyers navigate their options, this guide provides a framework for evaluating agentic AI for compliance, covering:

- Key differences among agentic AI solutions on the market
- What evaluation criteria matter most for success
- A scorecard for comparing potential vendors

Implementing agentic AI turns compliance from an inefficient requirement into a strategic lever. These systems automate investigative work, empowering compliance teams to focus on genuine risk— enabling growth without compromise.

Understanding Key Differences Between Agentic AI Solutions

Agentic AI solutions differ in how they're built, what they can do, and how they fit into your existing infrastructure. Before evaluating specific vendors, it's worth understanding the key distinctions between solution types.

Bank-Grade vs. Fintech-Grade AI

The first distinction is regulatory rigor: Does the solution have the auditability, testing, and controls that regulators expect, or will you need to build that infrastructure yourself?

Bank-grade agentic AI solutions are purpose-built to perform in regulated environments, with auditability and controls designed from the ground up. Fintech-grade solutions leverage general-purpose AI technology prioritizing speed to deployment and ease of integration. Where a solution falls on this spectrum affects the tradeoff between auditability and regulatory readiness and ease of integration and usability.

Bank-grade solutions will typically offer:

- Full audit trails with cited sources and traceable reasoning for every decision
- Rigorous model validation pre-launch, with ongoing monitoring for errors and drift post-launch
- Transparent testing methodologies and visibility into prompt performance over time

Fintech-grade solutions are characterized by:

- Proven, fintech-grade AI models applied to compliance workflows
- Seamless API integration with modern tech stacks
- Emphasis on ease of use and rapid deployment

Key consideration:

Bank-grade solutions are built to meet stringent regulatory requirements and auditability standards, but may require more complex integration. Fintech-grade solutions prioritize user experience and API compatibility but may leave some compliance gaps for banking use cases.

AI-Native vs. Bolt-On Automation

The second distinction is architecture: AI-Native solutions are standalone systems that layer on top of existing infrastructure via API, while bolt-on solutions embed AI features directly into platforms you may already use.

AI-native solutions offer:

- Compatibility with existing case management or screening infrastructure
- Advanced automation capabilities across complex workflows like enhanced due diligence and transaction monitoring
- Purpose-built AI architecture designed for flexibility across use cases

Bolt-on solutions are characterized by:

- Easy adoption if you're already using the core infrastructure
- Simple AI automations configured in a unified system
- AI capabilities defined by the host system's roadmap

Key consideration:

AI-native solutions offer greater flexibility, more advanced automation capabilities, and compatibility with existing case management systems but may require more complex integration work upfront. Bolt-on solutions simplify adoption but tend to offer simpler automations, less flexibility, and may create vendor lock-in.

AI Platform vs. AI Point Solutions

The third critical distinction is one of scope. Is it a point solution that automates a single compliance workflow, or is it a comprehensive platform solution that supports multiple compliance functions within a unified AI system?

Point solutions provide:

- Deep functionality for one specific workflow, such as sanctions screening or transaction monitoring
- Rapid time to value in the area of focus
- Simple implementation with minimal configuration required

Platform solutions enable:

- Broad coverage across sanctions, transaction monitoring, enhanced due diligence, periodic reviews, KYB, and more
- Expansion through configuration rather than new implementations
- API connectivity to existing case management and screening systems that connect data across workflows

Key consideration:

Point solutions deliver fast results for a well-defined workflow but may require multiple vendors as your automation needs expand. Platform solutions consolidate vendor management and connect data across compliance functions but require more configuration upfront.

Key Evaluation Criteria

Once you understand the solution landscape, it's crucial to evaluate specific solutions for the capabilities that matter most for project success. The following five key criteria will help identify solutions that will deliver fast time-to-value, operational control, and sustainable performance over time.

Integration Approach

What to look for: Fast implementation without infrastructure disruption

Solutions that layer on top of existing systems allow institutions to deploy AI capabilities quickly without ripping out what already works. Compliance teams will benefit from accelerated time-to-value and reduced implementation risk.

Best-in-class features:

- API connections to case management and screening systems
- Implementation timelines measured in weeks rather than months
- Minimal internal engineering resources required
- Preservation of validated Standard Operating Procedures

Meso deployed Greenlite in under two weeks with zero engineering resources. Learn how.

Read the case study ↗

Deployment Flexibility

What to look for: The ability to match automation levels to workflow risk profiles with progressive expansion

Compliance teams need control over how much autonomy they give AI systems. Phasing deployment strategically allows institutions to start with human oversight on high-risk workflows and expand automation as solution confidence builds.

Discover how HitPay scaled from KYB to EDD, sanctions, and periodic reviews—all on a single platform.

Explore the case study ↗

Best-in-class features:

- Human-in-loop and auto-close modes with control over whether AI outputs require human approval
- Risk-based queue sequencing to start with lower-risk alert types before expanding to higher-risk ones
- Workflow expansion from initial use case to additional compliance functions over time
- Pilot-to-production path with validation on historical alerts before deploying to live workflows

System Integrity

What to look for: The ability to produce accurate, explainable, and auditable decisions that can withstand regulatory scrutiny

Defensible AI decisions require traceable reasoning, cited sources, and quality control that meet or exceed what human analysts deliver, with full auditability. This addresses what regulators care about most: consistent quality, investigation depth, and rigorous documentation.

Best-in-class features:

- Traceable reasoning with cited sources for every decision
- Specific attributable factors (name similarity, date matches, geographic proximity) rather than confidence scores alone
- Auditable entity resolution, research steps, and scoring logic
- Quality control processes that uphold rigorous frameworks, including random sampling, notes review, and periodic assessment

Want to know how Fluz reduced false positives by 98% while raising compliance quality?

Check out the case study ↗

Regulatory Validation Framework

What to look for: Systematic testing and monitoring that supports model risk management requirements

Regulators expect institutions to validate AI systems with the same rigor they apply to any model or process. Strong validation frameworks enable institutions to apply existing QA processes without creating parallel structures, while providing the transparency needed for regulatory conversations.

Best-in-class features:

- Pre-launch validation through technical testing, expert compliance review, and accuracy benchmarks
- Prompt libraries with human-validated test sets
- Post-launch monitoring with error tracking, drift detection, and anomaly alerts
- Visibility into testing methodologies and prompt performance over time
- Regression testing when foundation models change

Grasshopper Bank cut due diligence time by 70% with AI-powered processes that earned positive OCC feedback. Explore how.

Read the case study ↗

Ongoing Maintenance Requirements

What to look for: Minimal post-deployment operational burden managed by the vendor

AI systems evolve continuously with new models, new workflows, updated integrations. Some solutions include partner-owned maintenance, enabling compliance and product teams to focus on complex investigations and high-priority operations instead.

Best-in-class features:

- Testing new foundation models against prompt benchmarks before production deployment
- Ongoing assessment of where AI delivers leverage versus where deterministic approaches remain appropriate
- Building and validating automation for additional workflows as institutional needs expand
- Bug fixes, integration updates, and performance optimization as systems evolve

Solution Evaluation Scorecard

Evaluating multiple solutions across this many dimensions can get unwieldy. This scorecard consolidates the key criteria into a side-by-side format, making it easier to compare vendors and identify the right fit for your compliance operations.

Criteria	Bretton AI	Partner 2	Partner 3
Bank-Grade vs. Fintech-Grade AI	Bank-Grade		
Platform vs. Point Solutions	Platform		
L1 Automation vs. Case Management	L1 Automation		
Fast implementation	✓		
Flexible deployment	✓		
System integrity	✓		
Regulatory validation	✓		
Partner-managed maintenance	✓		

Matching Solutions to Institutional Goals

Every institution has different priorities when selecting an agentic AI solution for compliance. You might need to deploy fast, prioritize regulatory defensibility, or plan for long-term scalability across multiple compliance functions. But regardless of where you start, the solution you choose should ultimately produce measurable business outcomes: reduced manual effort, faster case resolution, and demonstrable ROI.

Bretton AI's agentic AI is purpose-built to deliver these outcomes from day one. As a bank-grade, AI-native solution, Bretton AI layers on top of your existing case management and screening infrastructure to automate L1 investigative work across sanctions, transaction monitoring, enhanced due diligence, periodic reviews, and KYB.

Here's what that looks like in practice:

Fast time-to-value:
API-first integration enables production deployment in under four weeks, with minimal demand on your internal engineering resources.

Regulatory confidence:
Architecture designed for institutions regulated by the OCC, FDIC, and Federal Reserve delivers traceable reasoning, human-validated prompt libraries, and dedicated AI operations teams to support examination readiness.

Long-term scalability:

Coverage across 50+ use cases like sanctions, transaction monitoring, EDD, and more—with entity-level insights that connect investigations—enables you to expand automation without adding vendors.

Bretton AI customers have seen firsthand how these features work together to produce transformative results, like:

- 900% more screening alerts handled per month
- 58% reduction in analyst time on EDD cases
- 70% faster due diligence completion
- $5-10M+ in annual operational savings
- Preservation of validated Standard Operating Procedures

Bretton AI delivers the audit-ready, regulator-aligned AI infrastructure that compliance teams need to scale with confidence. See what's possible for your institution.

Book a demo today to see how Bretton AI can transform your compliance operations.