Put AI Agents Safely to Work in Your Bank with “AI GRCC”

Alec Crawford
Founder & CEO, Verapath
June, 2026

AI, Risk Management & Regulatory Expertise

- Building neural networks from scratch at Harvard
- CRO and running technology at Ziff Brothers Investments
- CRO, part of compliance, and AI as a Partner at Lord Abbett
- Independent Director of public company and Chair of Risk Committee
- Book: "Managing Global Risks and Opportunities in the AI Era", Macmillan
- Plus, banking!

Where Are You on Your AI Journey?

AI will change the world. We accelerate high-risk AI adoption by leveraging the best AI models, aligning AI with strategy and people, and making it safe, trustworthy, and compliant.

– Alec Crawford, Founder & CEO, Verapath

AI Philosophy – Stay in Control of AI and Your Data

- **Flexible AI**
  Works with any AI model, so you're not dependent on a single provider.
- **Own Your Data**
  Your data stays fully internal and never leaves your environment.
- **Easy Integration**
  Connects seamlessly with the tools you already use.
- **Smart Automation**
  A unified platform with AI assistants handling complex workflows.
- **Replace Your Apps**
  AI will replace apps, like your CRM, over time

Lead Your AI Transformation from the Inside Out

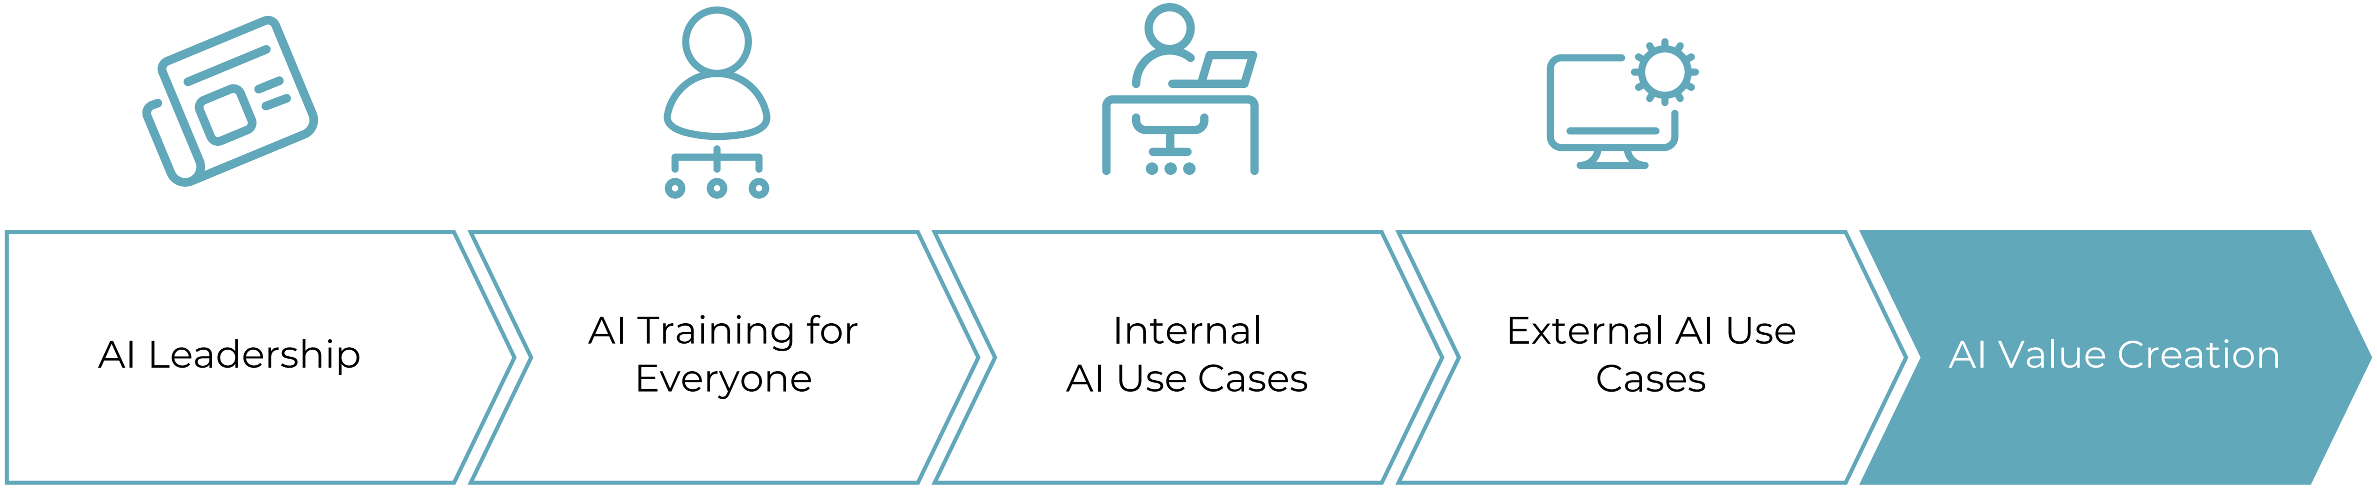

- **Align AI with Strategy**
- AI-aware leadership
- Appoint an AI champion

- **Mandatory training**
- Team-specific training
- Identify needs

- **Start with internal agents**
- Align with each team
- Review for compliance

- **Limited scope**
- Custom models
- Guard rails

- **Present KPIs**
- Efficiency gains
- FTE equivalents

Source: Verapath

Transform Your Bank with AI Agents

1. Low Hanging Fruit

- Easy AI agents like meeting summaries
- Personalized AI, e.g. for email
- Each team gets their own specific agents

2. Custom Workflows

- Customize new agentic workflows
- Connect to any model or data
- Add new agents/workflows every month

3. Digital Workers

- Digital workers performing complex tasks like loan closing
- Human in the loop
- Full audit trails and observability of AI agents and user interaction

Source: Verapath

How Do I Accelerate AI Adoption for Everyone?

Problem	Solution
AI Adoption: *How can we get 80%+ adoption?*	Customized AI training, change management, and use cases.
AI Data Integration: *How do we safely integrate AI across all our data and applications?*	Use APIs, MCP, and maintain ”permission awareness” for safety.
AI Control/Monitoring: *How do we control and monitor people using AI tools everywhere?*	Take control of AI by routing everything for safety, security, and compliance through an “AI hub”.
Getting <u>All</u> the AI Models: *How can I use one platform to manage access to all the AI models?*	Deploy AI models inside your private cloud and run through a hub like Verapath.

Source: Verapath

Solve the Hard AI problems and Create ROI

Deploy AI Safely for Everyone

The only platform combining Governance, Risk, Regulatory Compliance, and Cybersecurity in a single secure hub.

Automate Key Processes

Drive efficiency with agentic AI workflows and process management specifically designed for banking and financial operations.

Integrate Seamlessly

Connect with anything securely : e.g. Salesforce, core databases, document stores.

Use Any AI Model

Single interface for ChatGPT, Claude, Gemini, and more.

Real-Time Risk Management

Low-latency inspection and sensitive data protection.

Built-In Regulatory Compliance

Comply with AI regulations in finance from the SEC, FDIC etc. NIST and CRI compliance built in

AI Lifecycle Governance

End-to-end governance from use-case intake to audit trails.

Continuous Innovation

Monthly updates with the latest models and capabilities.

The New Approach Is Here

The FDIC and Cyber Risk Institute Have 230 Ideas for You

Problem
Governance: *Customers need to centrally control AI models, users and data.*
Risk Management: *Firms must protect proprietary and customer data by law*
Compliance: *AI compliance rules are complex and expanding*
Cybersecurity: *AI is the new attack surface for cybercriminals*

Source: Verapath

The CRI published their Financial Services Artificial Intelligence Risk Management Framework (FS AI RMF) in February, 2026 and It is becoming the de facto standard for banks in the US.

Common AI Agent Use Cases Today

There is a widening gap in finance

General	Human Resources	Risk	Compliance	Lending	CFO’s Office	Media and Comms
Secure Board meeting minutes	**Customized resume screener**	**Fraud detection**	**Firmwide policy Q&A chatbot**	**Automated credit memo drafting**	**Earnings call script writer**	**Customer Outreach**
Identifying and categorizing customer complaints	New “people manager” AI coach	Regulatory research/update	HMDA reporting via third-party API	Loan closing	Vendor contract analysis	Regulatory review for external messaging
RFP answers	HR policy chatbot & knowledge center	Portfolio risk monitoring	SAR reporting assistance	Data democratization	Bank competitor research	Social media content creation

Source: Verapath

And Near-Future Use Cases

AI will make finance more dangerous at first with "agentic commerce"

Potential Use Cases	Minimum Requirements
Shopping: *Instead of clicking on a website, a consumer asks their agent to buy something from sneakers to a plane ticket.*	Security, "know your agent", discovery of merchandise sellers, payment method (e.g. CC or stablecoin), guarantee framework, reg E compliance, consumer digital vault.
Investing: *Instead of a brokerage app, an AI agent could seek the best price across multiple platforms and other AI agents.*	Security, "know your agent", vetted investment platforms/agents, payment method, embedded trading algorithms, liability
Smart contracts: *Contracts on the blockchain (e.g. Ethereum) that automatically pay or transact when certain conditions are met.*	Contract creation and cybersecurity, tools to shift dollars into stablecoins or other crypto, KYA for AML purposes, legal enforceability, "oracle" for tamper-proof condition monitoring.

Source: Verapath

Who is going to provide the agents, the security layer, the "know your agent" framework, and the digital vault consumers will need to store their most private data securely?

AI Workflow Governance Is Hard!

Problem	Solution
AI Accuracy: *How can we get close to 100% accuracy in critical processes, such as lending?*	**Occam's Razor:** Use the simplest solution at each step, saving AI for when it is needed.
AI Data Integration*: How do we safely integrate AI across all our internal and external data, documents, emails, and third-party applications?*	**Own Your Data:** Keep it internal whenever possible, Integrate other data securely whether a database, documents, API, MCP, internal or external and maintain "permission awareness" for safety.
AI Control/Monitoring*: How do we control and monitor multiple AI tools across the organization for safety, security, and compliance?*	**Central AI Hub:** Take control of AI by routing people, permissions, data, AI agents, actions and models through an AI Hub. Full audit trail.
Using Many AI Models: *How can I use one platform to manage access to Open AI, Anthropic, Gemini, and the open-source models safely?*	**Private Model Deployment:** Deploy standard and custom AI model deployment inside your private cloud, on premises, or through API or MCP connections connected to your AI hub.

Source: Verapath

Case Study 1: Knowledge Center

50%+ overall savings

AI Knowledge Center

- **Reduced average call lengths 29%** from 8.5 minutes to 6 minutes
- Used the AI knowledge center to turn new call center employees into experts.
- Improved employee satisfaction in a high-turnover role

AI Customer Call Analysis

- **Reduced number of calls by 25%** by answering top 3 customer questions on their website right above call center phone number
- They processed 150,000 call transcripts and analyzed them in minutes.
- Identified top customer problems, best employees, product opportunities

This financial institution says they paid for their entire investment with this single use case. Knowledge centers can be used broadly for all employees for many purposes.

Case Study 2: AI-Powered Outreach

Create new capabilities with customized AI agents

Automating Analysis

Analyze customer behavior, bank engagement and statistics – for example:

Large checking account balance.

Creating Proposals

Create customized proposals and reach out to the customer with information:

A bonus CD rate.

Marketing Tools

Communicate the information to customers in the best way:

Send a personalized offer.

Customized AI agents can take your business to the next level in strategy, sales, marketing and more.

Case Study 3: Fighting Fraud

The Number 1 Problem for Banks Today

Automated Fraud Analysis

- **Scan incoming messages for signs of fraud.**
- Flag suspicious emails and customer behavior to employees, then customize.
- Rank “customer” behavior for risk, for example a suspiciously large wire transfer overseas..

Fraud Investigation

- **Fraud investigation AI agent and copilot.**
- Ask and get answers to questions about customer activity.
- Automate regulatory reports.

Each bank’s custom agents are proprietary and stay within their own firewall.

Custom Agentic AI Workflow Examples

Bank Individual Customer Onboarding QA

1. Missing/Expired Document Check
2. OFAC Check
PEP Check
Newly Banked
3. Stripe Account Match and Source of Funds
4. Review Credit Report and AML Risk Score
5. Complete

SBA Loan Prep for Approval Process

1. Entity & Personal Confirmation and Ownership Percent
2. Document Intake & Normalization
3. Cashflow, DSCR & Repayment Analysis
4. Risk Review: Bankruptcy, Criminal, Fraud
5. Forward to Human Underwriter

Source: AI Risk, Inc., additional one-time charge for Verapath to build custom agentic AI workflows

Lack of Risk Management?

Four Steps to Safely Manage AI across Your Organization

1. **Enterprise AI Platform:** Connect the best AI models, your data, and your teams together. Use AI workflows and find an AI agent library and tools <u>specifically</u> for finance.
2. **Real-Time AI Risk Management:** Risk manage all your production AI usage from users to models to AI agents, including low-latency prompt/response inspection, constraints, sensitive data protection, analytics, and real-time risk detection that <u>includes</u> third-party AI solutions.
3. **AI Lifecycle Governance:** Governance of AI initiatives across their full lifecycle, including use-case intake, ownership, risk mitigation, policy enforcement, model & vendor inventory, regulatory compliance management (US & Europe), and a complete immutable audit trail.
4. **Continuous Innovation:** Test and onboard the latest Gen AI models, add new AI agents and workflows, and roll out new capabilities.

Verapath

Secure: Run Private AI Inside Your Environment

Connecting data, people & models — securely

Trust your ai and control your data

AI BACKENDS

AI BACKEND	AI BACKEND	AI BACKEND	AI BACKEND	AI BACKEND	AI BACKEND
Azure Foundry	AWS	OpenAI	Anthropic	Google AI	Meta

EXPERIENCE

Web · Desktop · iOS · Android

One conversation surface across every device

CEO/COO · Strategic AI partner · agentic workflows

SECURITY

Your Internal Infrastructure Azure AWS

Control your data without worrying about a SAS Provider

CISO/CIO · Strategic AI partner · agentic workflows

OPERATIONS

Admin Console & Telemetry

Tenants, models, costs, usage — one pane of glass

CIO · Driving ROI with AI · one platform

COMPLIANCE

Audit, Archive & Reporting

Every prompt, response and tool call signed and retained

CCO · Regulatory compliance · expert advice

GreenLight Policy

OAuth2 · SSO

Verapath

AI GATEWAY

powered by Max

Workflows · Agents

Audit Trail

VERAPATH MCP RELAY

MCP Relay

Brokered tools, OAuth-scoped, fully logged

· AI cybersecurity · secure MCP hub

VERAPATH SCANNER

SCANNER

Brokered tools, OAuth-scoped, fully logged

· AI cybersecurity · secure MCP hub

3RD-PARTY APIS

Internal Databases

SQL, Mongo, Cosmos, Postgress and more

CTO · Connect to databases with any model

3RD-PARTY APIS

Any Third Party API

Automatic API importing from Swagger & Postman

CAIO · Fast easy no code integrations

PUBLIC

Internet & Web Search

Egress through policy + content scan

TRUSTED DATA CONNECTIONS

MS GRAPH

Outlook · Teams · SharePoint

Mail, chat, files via delegated OAuth

CEO/COO · Strategic AI partner · agentic workflows

INTERNAL

SQL · Files · Cosmos

Tenant data — never leaves the tenant

CDO · AI connected to anything · data experts

VERAPATH.SCANNER

Scanner Service

OCR, classification & PII scan

DATA & SYSTEMS

Compliant: Regulators Are Focused on AI in 2026

Compliance with the 230 CRI AI RMF requirements for financials Is built into Verapath

Governing
- Ingests policies & procedures
- Governs AI & data access
- Executes policy enforcement

Mapping
- Assigns AI agents & models
- Maps data connections
- Classifies AI risks

Measuring
- Creates auditable AI inventory
- Risk/red flag reporting
- Record use & feedback

Managing
- Human approvals
- Complex workflows
- Model/agent change controls

The CRI AI Risk Management Framework is emerging as a leading AI governance standard for banks and financial institutions. Developed with more than 100 financial institutions and Treasury and FDIC input, it translates the NIST AI RMF into 230 financial-services-specific control objectives.

Guard Rails for AI Inputs and Outputs

Financial Advisor	Compliance	General	Coding	Marketing	Data
Accidental customer information reveal	PII/PHI blocking or encryption	Hacking attempts	Code for non-IT users	Confidential data blocking	Strict data governance
Unsanctioned tax or legal advice	AI-assisted document or credential creation	Topic blocks: games, medical diagnosis, etc.	Revealing private keys/data	Reputational risk	PII/PHI redaction
Misuse of client data	Bias/toxic language	Competitor ban	Token limit enforcement	Misalignment with brand	File upload restrictions
Restrict which client data AI has access to	Avoid investment guarantee language	Capture and check invisible text	Unsanctioned code	Content scope	Screening for data poisoning

Source: AI Risk, Inc.

The Future of Money: Agentic Commerce

Huge first-mover advantages for banks in agentic commerce

Potential Use Cases	Minimum Requirements
Shopping: *Instead of clicking on a website, a consumer asks their agent to buy something from sneakers to a plane ticket.*	Security, ”know your agent”, discovery of merchandise sellers, payment method (e.g. CC or stablecoin), guarantee framework, reg E compliance, consumer digital vault.
Investing: *Instead of a brokerage app, an AI agent could seek the best price across multiple platforms and other AI agents.*	Security, ”know your agent”, vetted investment platforms/agents, payment method, embedded trading algorithms, liability
Smart contracts: *Contracts on the blockchain (e.g. Ethereum) that automatically pay or transact when certain conditions are met.*	Contract creation and cybersecurity, tools to shift dollars into stablecoins or other crypto, KYA for AML purposes, legal enforceability, “oracle” for tamper-proof condition monitoring.

Source: Verapath

Who is going to provide the agents, the security layer, the “know your agent” framework, and the digital vault consumers will need to store their most private data securely?

Champion Transformation as a Visionary AI Leader

- AI agents for each and every one

  Increase efficiency across your company on day one. Start with dozens of our built-in AI agents for each department and expand with custom agents while preserving data security.

- A platform, not just an app

  AIR is a platform for Gen AI – a holistic solution for companies to use all kinds of Gen AI safely and securely across your teams and use cases.

- Work is changing forever

  Secure, agentic AI workflow will change business forever. AI will change the value of different skills and roles in the near future. Bring it to your company and team.

- Beyond AI trust and safety

  Governance, risk management, regulatory compliance, and cybersecurity go way beyond the current category of "AI trust and safety". Large technology companies will not want to deal with industry-specific regulations.

- Must-have cybersecurity & compliance

  Dashboard and alerting 24/7 for management and cybersecurity teams showing usage, geography, compliance, hacking attempts, and summarizing activity.

- Future proof your organization

  Do not lock yourself into one AI model provider as we do not know who will be best, or cheapest, a year from now. AIR is compatible simultaneously with all proprietary and third-party AI LLM systems.

Resources

White Papers

- *A Risk Management Framework for Large Language Models* (April 2024)
- *How to Regulate AI for Banks* (June 2025)
- *From SEO to GEO: How AI Is Reshaping Search (and How to Take Advantage of It Today)* (September 2025)

Magazine Articles (in *Directors and Boards*)

- *The Board and AI's Human Capital Reset (June, 2024)*
- *Overcoming the Three Hidden Dangers of AI (March, 2025)*

Podcasts/Blog

- Substack blog *AI: Trust but Verify* and monthly AI Newsletter on *LinkedIn*
- Host of the top podcasts *AI: Trust but Verify and Future Tech Banker*

Courses

- Stamford has released all their AI courses for free!
- Cassie Kozyrkov (ex-Google) on YouTube and Andrew Ng on Coursera

Access the Verapath Content Vault Now:

Appendix

FAQ on Real-Time AI Risk Management

1. **Which types of AI interactions can your platform inspect?** The platform can inspect input text, output responses, OCR, file uploads as well as pulls from external sources such as document storage or databases.
2. **Is inspection real-time or near real-time?** Inspection is real-time with 500 ms typical lag time.
3. **What types of sensitive data can your platform detect?** The platform detects global PII, PCI, PHI, and financial data to ensure protection and compliance.
4. **How are sensitive data events handled?** Sensitive data may be blocked, redacted, or encrypted for compliance and security. Everything is logged.
5. **What automated or manual actions can be triggered when unapproved AI use or policy violations are identified?** Automated alerts and manual reviews are initiated upon detection of unapproved AI use or policy violations.
6. **What reporting dashboards are available for usage analytics, risk events, and compliance tracking?** Comprehensive dashboards, ticketing systems, reports and an SEC-compliant e-discovery tool provide insights into usage analytics, risk events, and allow compliance tracking. Standard and customizable dashboards are available and may be stored for easy updates 24/7.
7. **What export or API integrations exist to feed analytics into enterprise systems?** The platform supports multiple export formats and API integrations for feeding analytics into enterprise systems including, but not limited to, JSON and streaming responses and secure MCP hub integration**.**
8. **How does your solution identify and mitigate harmful or noncompliant AI-generated content?** Noncompliant content is identified and mitigated through robust filtering systems and blocked. The system may be configured in detail, e.g., preventing AI medical advice at a bank.
9. **What mechanisms are in place to detect prompt injection or adversarial activity**? The system includes a database of prompt injections and other attacks on Gen AI. Detection mechanisms include real-time monitoring and advanced threat detection algorithms to counter prompt injections and adversarial activities.
10. **Do you support agentic AI monitoring (task chaining, decision logging)?** Agentic AI monitoring via task chaining and decision logging is supported as well as version control, a monitoring dashboard, governance and safety controls, and the integration layer. Observability control can be integrated into the system via third-party tools.

FAQ on AI Lifecycle Governance

1. **How does your solution capture and catalog AI use cases, models, and associated datasets across the organization?** The solution provides a centralized system to capture and catalog AI use cases, models, integrations, and datasets seamlessly as they are added or removed by the administrators.
2. **What metadata is tracked for each?** Metadata tracked includes owner, business purpose, data sensitivity, model type, and source and may be customized to include lifecycle stage or other information for the client.
3. **Can both internally developed and third-party/SaaS AI solutions be registered and governed?** Yes, the platform allows registration and governance of both internally developed and third-party/SaaS AI solutions including CoPilot, Microsoft Foundry, AWS Bedrock, the major GenAI model providers, and third-party AI tools with appropriate API or MCP access.
4. **How are governance workflows (e.g., intake, review, approval, decommissioning) managed within the platform?** There is a user interface on the platform specifically for administrators and a separate tool for compliance professionals. Agents and workflows go through an intake, review, approval, and decommissioning process within that platform.
5. **Can use case assessments be tailored to financial-services regulations (GLBA, SEC/FINRA, OCC, NIST)?** Yes, the platform was designed from the ground up to support financial-services regulations such as GLBA, SEC/FINRA, and the OCC as well as the NIST AI Risk Management Framework (RMF). We integrate with third-party tools for Fair Lending and other model evals.
6. **Can governance data or audit evidence be accessed via API for enterprise reporting?** Yes, governance data and audit evidence can be accessed via API to facilitate enterprise reporting and connections to other systems.
7. **What deployment models are supported (SaaS, private cloud, on-premises, hybrid)?** The solution supports private cloud, on-premises, and hybrid deployment models at different price points. We do not provide a SaaS solution for security and latency reasons.
8. **How does your solution scale across multiple entities, business units, and varied data environments?** The solution has flexibility and can scale on a private cloud to over 1 million users at dozens of sites simultaneously. It can connect to virtually any database, API, or MCP connection. In addition, it can be separately deployed or be partitioned inside the software to reflect different business entities.

FAQ on AI Regulatory Compliance

1. **What are regulators saying about using AI in financial services?** Regulatory bodies in the US have stated "just because it is AI, the existing rules still apply. The SEC's Reg S-P impacts AI. The FDIC partnered with the Cyber Risk Institute to release 240 specific guidelines for AI in February 2026, which aligns with the National Institute of Science and Technology (NIST) AI Risk Management Framework.
2. **What risk management frameworks are US regulators encouraging?** The NIST AI Risk Management Framework has been adopted by larger financial institutions in the US. This guideline requires AI policies and procedures to govern, map, measure, and manage your AI ecosystem, including the users. Part of this is an immutable database to track prompts and responses as well as AI agent and model creation, updates, and deprecation.
3. **What reporting dashboards are available for compliance tracking on the Verapath platform?** Comprehensive dashboards, ticketing systems, reports and a regulatory compliant e-discovery tool provide insights into usage analytics, risk events, and allow compliance tracking across hundreds of items. Compliance and regulatory searches are stored in their own immutable database.
4. **How does the Verapath platform help comply with financial-services regulations (GLBA, SEC/FINRA, OCC, NIST)? T**he platform was designed from the ground up to support financial-services regulations such as GLBA, SEC/FINRA, and the OCC as well as the NIST AI Risk Management Framework (RMF). There are tools for everything from PII redaction to topic blocking.
5. **Can compliance data or audit evidence be accessed via API for enterprise reporting?** Yes, governance data and audit evidence can be accessed via API to facilitate enterprise reporting and connections to other systems.
6. **What are regulatory expectations on supervision when AI is used in communication or investing tools?** AI outputs must be supervised for accuracy, fairness, suitability and conflicts (in general UDAAP-type issues). Books and record rules also apply, including chats, investment research, emails, and other communications. Verapath allows for customized retention policies.
7. **What disclosures should we provide to clients when AI is used in advice, personalization, or processes?** Client disclosures around transcription, material AI use for communications, advice, or investing are required. It's also important not to "AI-wash" or overpromise on what AI can deliver. Tailor the disclosure to the use case: advice, personalization, and optimization, for example.