elvex

The Complete Security & Compliance Checklist for AI in Financial Services

Updated November 2026

elvex.ai

Secure By Design

Financial institutions face a critical paradox: AI promises transformational efficiency gains, but regulatory requirements and security concerns create seemingly insurmountable barriers to adoption. The question isn't whether to adopt AI, but how to do it safely, securely, and in full compliance with the regulatory frameworks that govern your industry.

This guide provides a comprehensive security and compliance framework for deploying AI in financial services, addressing FINRA, SOX, GLBA, and other regulatory requirements specific to regulated institutions.

In this guide:

- The Security Foundation: What Actually Matters
- Architecture for Security: Building in Protection from Day One
- Regulatory Compliance: Meeting Industry-Specific Requirements
- Mitigating AI-Specific Risks
- Governance and Control at Scale
- The Complete Checklist
- Entry-Level SOX Playbook for Financial Services

For more free resources:

1. Security Questions to Ask AI Vendors: A Financial Services Guide
2. More on AI in financial services on the elvex blog

The Security Foundation

Certification Standards as Your Starting Point

SOC 2 Type 2 certification establishes the baseline security posture for any AI platform operating in financial services. SOC 2 Type 2 signals that a vendor has implemented and maintained critical security controls over time, including:

- **Principle of least access:** Users only have access to the data and systems they need for their specific role
- **Encryption everywhere:** Data encrypted both in transit and at rest, with no exceptions
- **Quarterly access reviews:** Regular audits of who has access to what, with documented justification.

- **Sensitive data handling protocols:** Clear policies for who can touch or view customer information

Beyond SOC 2 Type 2, more stringent frameworks like ISO 27001 or HIPAA certification may be required depending on your specific use cases and regulatory environment.

> **Key Question to Ask Vendors** What security framework have you certified against, and can you provide evidence of ongoing compliance monitoring?

The Data Exfiltration Problem: And How to Solve It

When AI platforms first emerged, the primary concern from enterprises was simple: our proprietary data will be used to train models, and our competitive secrets will leak.

Early solutions attempted to solve this through obfuscation, removing sensitive information before sending it to models. This approach fails because:

- It's a whack-a-mole game trying to identify every possible sensitive data point
- Determining what's "sensitive" often requires an LLM itself
- Even with filtering, you can't guarantee complete protection

The Modern Solution: Self-hosted models eliminate the training data concern entirely. Today's options include:

- **AWS Bedrock:** Host models within your AWS environment
- **Google Vertex AI:** Deploy models in your Google Cloud infrastructure
- **Azure OpenAI:** Run OpenAI models in your Azure tenant

With these approaches, model providers contractually guarantee your data won't be used for training purposes, and you maintain full control over where your data lives.

> **Key Question to Ask Vendors** Does your platform support self-hosted model providers, and can I switch between them without rebuilding my workflows?

On-Premises vs. Cloud: Making the Right Choice

For most financial services use cases, a cloud-based platform with proper security controls (SOC 2 Type 2, encryption, access controls) provides sufficient protection. The question of on-premises deployment typically arises only for:

- Extremely sensitive use cases involving classified information

- Regulatory requirements specific to certain jurisdictions
- Large-scale deployments where the economics justify the complexity

The Trade-offs of On-Premises

- Significantly higher implementation costs (typically seven-figure deals)
- Complex update and maintenance cycles (quarterly updates vs. continuous deployment)
- Reduced agility in adopting new features and capabilities

The Pragmatic Approach: Start with a cloud-based deployment for most use cases, while restricting certain highly sensitive data from the platform until you can justify an on-premises deployment.

> **Key Question to Ask Vendors** What's your approach to on-premises deployment, and at what scale does it become economically viable?

One Platform. Unlimited Possibilities. Any model. Any tool. Any employee. Completely secure.

Architecture for Security

Logical vs. Physical Data Separation

Understanding how your AI platform separates customer data is critical for assessing security risk.

- **Physical Separation:** Each customer gets their own database instance, making it physically impossible for data to leak between tenants. This approach is expensive and typically reserved for the most sensitive deployments.
- **Logical Separation:** All customers share database infrastructure, but data is separated through application-level controls (company IDs, access control lists, etc.). This is the standard approach for modern SaaS platforms.

Logical separation is not inherently less secure than physical separation, provided the platform:

- Implements proper access controls at the application layer
- Undergoes regular penetration testing to identify potential cross-tenant vulnerabilities
- Maintains comprehensive audit logs of all data access

> **Key Question to Ask Vendors** Do you use logical or physical data separation, and what penetration testing do you conduct to validate security?

Role-Based Access Controls (RBAC)

Security doesn't end at the platform level. The AI tools your employees use must have granular permission structures built in from day one.

Essential RBAC Features

- **User roles with distinct permissions:** Admins, builders, and end-users should have different capabilities
- **Resource-level sharing controls:** Data sources, assistants, and workflows should have configurable sharing (private, team, company-wide)
- **Group-based provisioning:** Automatically assign users to groups based on their role or department
- **SAML/SSO integration:** Sync permissions with your existing identity provider

The benefit of RBAC is that it scales with your organization. As you add more users, teams, and use cases, you don't need to manually manage permissions for every individual.

> **Key Question to Ask Vendors** How do you implement role-based access controls, and can you integrate with our existing identity provider?

Decoupled Data Sources

Many AI platforms tightly couple data sources with specific workflows or assistants. This creates a management nightmare as you scale.

The Problem with Coupled Architectures:

- Every new workflow requires re-uploading or reconnecting data
- No centralized view of what data has been connected to AI systems
- Difficult to audit which workflows have access to sensitive information
- Impossible to quickly revoke access when someone leaves or changes roles

The Solution: Decoupled Data Sources

- Data sources exist as independent, reusable components
- Connect a data source once, use it across multiple assistants and workflows
- Centralized access control and audit trail for all data
- Easy to see which workflows are using which data sources

This architectural decision has massive implications for security and compliance. With centralized data sources, you can:

- Quickly audit all AI workflows using customer PII
- Revoke access to a data source and immediately affect all dependent workflows
- Maintain a single source of truth for data governance

> **Key Question to Ask Vendors** Are data sources independent components, or are they tied to specific workflows?

Regulatory Compliance: Meeting Industry-Specific Requirements

FINRA Compliance for AI Systems

FINRA has made clear that its rules are technology-neutral and apply to AI just as they apply to any other tool. FINRA Regulatory Notice 24-09 specifically addresses generative AI and LLMs.

Supervision (FINRA Rule 3110)

Member firms must have reasonably designed supervisory systems tailored to their business. If using AI tools, policies and procedures must address:

- Technology governance
- Model risk management
- Data privacy and integrity
- Reliability and accuracy of AI models

Record-Keeping

All AI-generated communications, recommendations, and decisions must be retained according to FINRA record-keeping requirements. This includes:

- Customer communications generated or assisted by AI
- AI-generated research or recommendations
- Audit trails of AI system decisions

Content Standards

AI-generated communications must comply with FINRA's content standards, including:

- Fair and balanced presentation
- No false or misleading statements

- Appropriate disclosures and risk warnings
- Compliance with Regulation Best Interest (Reg BI)

Human-in-the-Loop

While not explicitly required by FINRA, supervisory obligations effectively mandate human oversight of AI outputs, particularly for:

- Investment recommendations
- Customer communications
- Suitability determinations
- Compliance decisions

> **Key Question to Ask Vendors** How does your platform support FINRA recordkeeping requirements and enable human oversight of AI-generated outputs?

SOX Compliance and IT General Controls (ITGCs)

The Sarbanes-Oxley Act requires controls around financial reporting systems. When AI becomes part of your financial reporting infrastructure, SOX IT General Controls apply.

Essential SOX ITGCs for AI Systems

Access Controls:

- User access management with documented approval processes
- Privileged access restricted to authorized personnel only
- Regular access reviews (quarterly or more frequent)
- Immediate revocation of access when employees leave or change roles

Change Management:

- Documented change control procedures for AI models and workflows
- Testing and approval requirements before production deployment
- Rollback procedures for failed changes
- Audit trail of all changes to production systems

Segregation of Duties:

- No single person should control the entire AI system lifecycle
- Separate roles for development, testing, and production deployment
- Independent validation of models before production use

- Separation between those who build workflows and those who approve them

Data Security and Integrity:

- Encryption of financial data at rest and in transit
- Backup and recovery procedures for AI systems
- Data validation controls to ensure accuracy
- Protection against unauthorized modification

Monitoring and Logging:

- Comprehensive audit logs of all system activity
- Automated monitoring for suspicious activities
- Regular review of logs by compliance personnel
- Retention of logs per regulatory requirements

> **Key Question to Ask Vendors** What SOX-compliant controls are built into your platform, and how do you support segregation of duties?

GLBA Safeguards Rule and NPI Protection

The Gramm-Leach-Bliley Act requires financial institutions to protect customer Nonpublic Personal Information (NPI). When AI systems process NPI, GLBA compliance becomes critical.

GLBA Safeguards Rule Requirements

Written Information Security Plan:

- Risk assessment identifying threats to NPI
- Safeguards designed to control identified risks
- Regular monitoring and testing of safeguards
- Oversight of service providers handling NPI

Encryption Requirements:

- NPI must be encrypted in transit and at rest
- Strong encryption standards (AES-256 or equivalent)
- Secure key management practices
- Encryption of data sent to AI models

Access Controls for NPI:

- Restrict access to NPI to authorized personnel only

- Multi-factor authentication for systems containing NPI
- Automatic logout after periods of inactivity
- Monitoring of access to NPI

Data Sanitization:

- Secure disposal of NPI when no longer needed
- Data minimization practices (collect only what's necessary)
- Automatic redaction of NPI in AI outputs when appropriate
- Procedures for de-identifying data for analytics

Third-Party Oversight:

- Due diligence when selecting AI vendors
- Contractual requirements for NPI protection
- Regular monitoring of vendor security practices
- Right to audit vendor controls

Incident Response:

- Written incident response plan for data breaches
- Notification procedures for affected customers
- Reporting to regulators as required
- Post-incident review and remediation

Example GLBA Privacy Notice Language for AI

> "We use artificial intelligence and machine learning technologies to improve our services and protect your account. These technologies may analyze your account information and transaction history to detect fraud, provide personalized recommendations, and enhance your experience. We maintain strict controls to protect your information, including encryption, access controls, and regular security testing. Your information is never used to train public AI models. You can opt out of certain AI-powered features by contacting us at [contact information]."

> **Key Question to Ask Vendors** How do you help financial institutions comply with GLBA Safeguards Rule requirements, and what controls protect NPI?

Mitigating AI-Specific Risks

The Hallucination and Prompt Injection Problem

AI models have inherent limitations that no platform can completely eliminate:

- **Hallucinations:** Models generating false or fabricated information. While less common with modern models, they still occur.
- **Prompt Injection Attacks:** Malicious instructions hidden in documents or web pages that cause the model to ignore its original instructions and follow attacker-controlled commands.

Why This Matters in Financial Services

- A hallucinated investment recommendation could lead to unsuitable advice
- A prompt injection attack could exfiltrate customer data
- Compliance violations could occur without the user's knowledge

Mitigation Strategies

1. **Human-in-the-loop for critical decisions:** AI should assist, not replace, human judgment for investment advice, underwriting decisions, and compliance determinations
2. **Approval workflows for actions:** Any action taken on behalf of a user (sending an email, updating a CRM, executing a trade) should require explicit human approval
3. **Comprehensive audit logs:** Every AI interaction should be logged with full context (user, data accessed, model used, output generated, actions taken)
4. **Scoped assistants with guardrails:** Build assistants for specific use cases with explicit instructions about what they should and shouldn't do
5. **Internal use cases first:** Deploy AI for employee productivity before customer-facing applications

> **Key Question to Ask Vendors** What controls do you have in place to prevent AI from taking unauthorized actions, and how comprehensive are your audit logs?

AI Disclosure Requirements

Financial institutions must be transparent about their use of AI, both to customers and regulators.

Customer Disclosure Best Practices

When to Disclose:

- When AI makes or significantly influences decisions affecting customers
- When AI is used in customer-facing applications (chatbots, robo-advisors)
- When AI processes customer data for personalization or recommendations
- In privacy policies and terms of service

What to Disclose:

- That AI is being used and for what purposes
- What data is being processed by AI systems
- How customers can opt out or request human review
- How to contact the institution with questions or concerns

Example AI Disclosure Language

For Terms of Service:

> "Our platform uses artificial intelligence and machine learning to enhance your experience, detect fraudulent activity, and provide personalized recommendations. AI systems may analyze your account information, transaction history, and interactions with our services. All AI systems are subject to human oversight, and you may request human review of any AI-assisted decision by contacting us at [contact information]."

For Investment Advisory Services:

> "We use AI-powered tools to assist our advisors in analyzing market data, identifying investment opportunities, and monitoring portfolio performance. All investment recommendations are reviewed and approved by licensed investment professionals before being presented to clients. AI tools do not make final investment decisions."

For Fraud Detection:

> "We use AI systems to monitor transactions and detect potentially fraudulent activity. If our systems flag unusual activity, we may contact you to verify the transaction or temporarily restrict account access. You can always contact us to discuss any fraud alerts or account restrictions."

Regulatory Disclosure Requirements

Financial institutions should be prepared to explain to regulators:

- Complete inventory of AI systems in use
- Purpose and use cases for each AI system
- Data sources and model providers
- Validation and testing procedures
- Governance and oversight structure
- Incident response procedures for AI failures

> **Key Question to Ask Vendors** What tools do you provide to help us maintain an AI inventory and generate regulatory disclosures?

Audit Trails, Data Retention, Reporting

Comprehensive logging and retention are fundamental to regulatory compliance in financial services.

Audit Trail Requirements

Every AI interaction should capture:

- **User information:** Who initiated the request
- **Timestamp:** When the interaction occurred
- **Data accessed:** What information was retrieved or processed
- **Model used:** Which AI model generated the output
- **Input and output:** The prompt and the response
- **Actions taken:** Any automated actions triggered by the AI
- **Human review:** Whether and when a human reviewed the output

Data Retention Requirements

Different regulations impose different retention periods:

Regulation	Retention period
FINRA	Most records must be retained for 3-6 years
SEC	Investment adviser records typically 5 years
SOX	Financial records for 7 years
GLBA	Customer information for duration of relationship plus applicable period

Your AI platform must support:

- Configurable retention periods by data type
- Immutable logs that cannot be altered
- Secure archival of historical data
- Easy retrieval for audits and investigations

Reporting Capabilities

Financial institutions need to generate reports for:

- **Internal audit:** Regular reviews of AI usage and compliance
- **Regulatory examinations:** Ad-hoc requests from regulators
- **Board reporting:** Summary of AI risks and incidents
- **Incident response:** Detailed forensics when issues occur

Essential reporting features:

- Searchable and filterable audit logs
- Export capabilities (CSV, PDF, API)
- Customizable report templates
- Scheduled automated reports
- Real-time dashboards for monitoring

> **Key Question to Ask Vendors** What are your data retention capabilities, and can you export comprehensive audit logs for regulatory examinations?

Governance and Control at Scale

Establishing AI Governance

Effective AI governance requires structure, accountability, and ongoing oversight.

AI Governance Committee:

- Reviewing and approving high-risk AI use cases
- Setting policies for acceptable AI use
- Monitoring AI-related risks and incidents
- Ensuring compliance with regulations
- Reporting to board and senior management

Committee Composition:

- Chief Risk Officer or equivalent
- Chief Information Security Officer
- Chief Compliance Officer
- Business line representatives
- Technology/AI specialists
- Legal counsel

Acceptable Use Policies:

- What AI systems are approved for use
- What data can be processed by AI
- What use cases require additional approval
- What outputs require human review
- How to handle AI errors or failures

Third-Party Vendor Oversight:

- Conduct due diligence before selection
- Review security and compliance certifications
- Negotiate appropriate contractual protections
- Monitor vendor performance and incidents
- Maintain right to audit
- Have exit strategy if vendor fails to perform

Board-Level Oversight:

- Receive annual briefing on AI program
- Review significant AI-related risks and incidents
- Approve AI governance policies
- Ensure adequate resources for AI risk management
- Hold management accountable for AI compliance

> **Key Question to Ask Vendors** What governance frameworks do you recommend, and how does your platform support centralized oversight?

Centralized AI Governance in Practice

The biggest security risk in financial services comes from the rogue AI usage happening across your organization right now, rather than from the AI platform you've vetted and approved.

The Rogue AI Problem:

- Advisors using ChatGPT for client research
- Analysts using Claude for investment memos
- Operations teams using Gemini for process automation
- Customer data and proprietary information leaking to public models

The Solution: Centralized AI Adoption

Rather than trying to prevent AI usage (which is impossible), provide a sanctioned platform that:

- Meets all your security and compliance requirements
- Is easier to use than consumer AI tools
- Integrates with the systems your employees already use
- Provides the capabilities they need for their specific roles

When you centralize AI adoption, you gain:

- Visibility into all AI usage across the organization
- Ability to enforce security and compliance policies
- Consolidated vendor management and cost control
- Consistent audit trails for regulatory examinations

> **Key Question to Ask Vendors** How do you help organizations consolidate AI usage onto a single, governed platform?

Cost Visibility and ROI Tracking

Security and compliance extend beyond preventing bad outcomes to demonstrating value and justifying investment.

Essential Analytics Capabilities:

- Which users are actively using AI
- Which assistants and workflows are driving the most value
- What models are being used and at what cost
- Time saved on specific processes (onboarding, underwriting, compliance)
- Adoption rates across departments and roles

This data serves multiple purposes:

- Justify AI investment to the board and regulators
- Identify high-value use cases to expand

- Find low-value use cases to eliminate
- Optimize model selection for cost and performance
- Demonstrate compliance with AI governance policies

> **Key Question to Ask Vendors** What analytics do you provide to track adoption, usage, and ROI?

The Complete Checklist

Use this checklist when evaluating AI platforms for financial services:

Security & Compliance Foundations

- [ ] SOC 2 Type 2 certified (or higher)
- [ ] Data encrypted in transit and at rest
- [ ] Support for self-hosted models (AWS Bedrock, Azure OpenAI, Google Vertex AI)
- [ ] Regular third-party penetration testing
- [ ] Contractual guarantees that data won't be used for model training
- [ ] Clear data residency options for regulatory requirements

Access Control & Permissions

- [ ] Role-based access controls (RBAC)
- [ ] SAML/SSO integration with your identity provider
- [ ] Group-based provisioning and management
- [ ] Granular sharing controls for data sources, assistants, and workflows
- [ ] Principle of least access enforced by default
- [ ] Quarterly access reviews supported

Architecture & Design

- [ ] Logical data separation with proven security controls
- [ ] Decoupled data source architecture
- [ ] Model-agnostic platform (switch providers without rebuilding)
- [ ] API access for custom integrations
- [ ] Scalable architecture for enterprise deployment

FINRA Compliance

- [ ] Comprehensive recordkeeping of all AI interactions
- [ ] Support for human-in-the-loop workflows
- [ ] Content standards enforcement for AI-generated communications
- [ ] Supervision and oversight capabilities
- [ ] Retention of AI-generated customer communications

SOX IT General Controls

- [ ] Documented access control procedures
- [ ] Change management with testing and approval workflows
- [ ] Segregation of duties between development and production
- [ ] Audit trails of all system changes
- [ ] Data backup and recovery procedures

GLBA Safeguards Rule

- [ ] Encryption of NPI in transit and at rest
- [ ] Access controls specific to NPI
- [ ] Data sanitization and minimization capabilities
- [ ] Automatic redaction of sensitive information
- [ ] Third-party vendor oversight support
- [ ] Incident response capabilities

AI-Specific Risk Mitigation

- [ ] Human-in-the-loop for critical decisions
- [ ] Approval workflows for actions taken on behalf of users
- [ ] Ability to set guardrails and constraints on assistants
- [ ] Support for internal use cases before customer-facing deployment
- [ ] Protection against prompt injection attacks

Audit, Retention & Reporting

- [ ] Immutable audit trails with full context

- [ ] Configurable retention periods by data type
- [ ] Ability to export logs to centralized logging systems
- [ ] Search and filter capabilities for investigations
- [ ] Customizable reporting templates
- [ ] Real-time monitoring dashboards

AI Disclosure & Transparency

- [ ] AI inventory and tracking capabilities
- [ ] Tools to generate customer disclosures
- [ ] Documentation for regulatory examinations
- [ ] Transparency into which models are used where

Governance & Control

- [ ] Centralized platform to consolidate AI usage
- [ ] Visibility into all AI interactions across the organization
- [ ] Support for AI governance committee workflows
- [ ] Policy enforcement capabilities
- [ ] Cost tracking and optimization tools
- [ ] Adoption and usage analytics
- [ ] ROI measurement capabilities

Vendor Considerations

- [ ] Transparent pricing without hidden AI premium fees
- [ ] Right to audit vendor controls
- [ ] Clear data ownership and portability terms

Entry-Level SOX Playbook for Financial Services

For organizations just beginning their AI journey, here's a practical starting point for SOX compliance:

Phase 1: Establish Foundation (Months 1-2)

Form AI Governance Committee:

- [ ] Identify committee members from risk, compliance, IT, and business
- [ ] Define charter and meeting cadence
- [ ] Establish escalation procedures

Document Current State:

- [ ] Inventory existing AI usage (including shadow AI)
- [ ] Identify financial reporting systems that use or could use AI
- [ ] Document data flows and access patterns

Define Acceptable Use Policies:

- [ ] What AI tools are approved
- [ ] What data can be processed
- [ ] What use cases require approval
- [ ] Training requirements for users

Phase 2: Implement Controls (Months 3-4)

Access Controls:

- [ ] Implement RBAC for AI platform
- [ ] Integrate with identity provider (SSO/SAML)
- [ ] Document access approval procedures
- [ ] Schedule quarterly access reviews

Change Management:

- [ ] Define change control procedures for AI workflows
- [ ] Establish testing requirements
- [ ] Create approval workflows
- [ ] Document rollback procedures

Segregation of Duties:

- [ ] Separate development and production environments
- [ ] Define roles (builder, reviewer, approver)
- [ ] Implement approval workflows
- [ ] Document role assignments

Phase 3: Monitor and Validate (Months 5-6)

Audit Logging:

- [ ] Enable comprehensive logging
- [ ] Configure log retention
- [ ] Set up monitoring alerts
- [ ] Test log export capabilities

Testing and Validation:

- [ ] Test access controls
- [ ] Verify change management procedures
- [ ] Validate segregation of duties
- [ ] Document test results

Board Reporting:

- [ ] Prepare summary of AI program
- [ ] Document risks and controls
- [ ] Present to board or audit committee
- [ ] Obtain formal approval

Phase 4: Continuous Improvement (Ongoing)

Regular Reviews:

- [ ] Quarterly access reviews
- [ ] Annual policy reviews
- [ ] Ongoing monitoring of AI usage
- [ ] Regular testing of controls

Adaptation:

- [ ] Update policies as regulations evolve
- [ ] Incorporate new use cases
- [ ] Refine controls based on lessons learned
- [ ] Expand to additional business areas

Key Questions to Ask Vendors

Below are the aggregated key questions to ask vendors when assessing AI platforms.

1. What security framework have you certified against, and can you provide evidence of ongoing compliance monitoring?
2. Does your platform support self-hosted model providers, and can I switch between them without rebuilding my workflows?
3. What's your approach to on-premises deployment, and at what scale does it become economically viable?
4. Do you use logical or physical data separation, and what penetration testing do you conduct to validate security?
5. How do you implement role-based access controls, and can you integrate with our existing identity provider?
6. How does your platform support FINRA recordkeeping requirements and enable human oversight of AI-generated outputs?
7. What SOX-compliant controls are built into your platform, and how do you support segregation of duties?
8. How do you help financial institutions comply with GLBA Safeguards Rule requirements, and what controls protect NPI?
9. What controls do you have in place to prevent AI from taking unauthorized actions, and how comprehensive are your audit logs?
10. What tools do you provide to help us maintain an AI inventory and generate regulatory disclosures?
11. What are your data retention capabilities, and can you export comprehensive audit logs for regulatory examinations?
12. What governance frameworks do you recommend, and how does your platform support centralized oversight?
13. How do you help organizations consolidate AI usage onto a single, governed platform?
14. What analytics do you provide to track adoption, usage, and ROI?

Security as Enabler, Not Barrier

The institutions that will win with AI are those that move fastest while maintaining security, compliance, and control.

The right approach to AI security in financial services centers on enabling safe, compliant, and auditable adoption at scale.

When you implement the framework outlined in this guide, you transform security from a barrier to an enabler. You can confidently deploy AI across your organization, knowing that:

- Your data is protected with financial-grade security
- Your compliance obligations are met across FINRA, SOX, and GLBA

- Your employees have the tools they need without resorting to rogue AI usage
- Your leadership has the visibility and ROI metrics they need to justify investment
- Your regulators can see clear evidence of responsible AI governance

The future of financial services is AI-native. The question is whether you'll get there securely, or whether you'll be forced to retrofit security onto rushed implementations.

elvex is the enterprise AI agent platform built for regulated industries. elvex gives financial institutions the freedom to build AI workflows with security, compliance, and audit trails built in from day one.

Talk to our team at elvex.ai/enterprise