Optiri
A Trellance Company

AI Security & Governance for Credit Unions

Managing Shadow AI, Embedded AI & Agentic AI Risks

Introductions

Shane Butcher
Executive Director, Optiri
CRO, Trellance

Barry Lewis
Senior Director, Security
and Technology Consulting

Today's Agenda

01 **AI Threat Landscape**
Shadow AI, Embedded AI, Agentic AI

02 **Credit Union Risk Profile**
Why you're high-value targets

03 **Regulatory Expectations**
NCUA, NIST, FFIEC

04 **Building Governance**
Teams, policies, technical controls

05 **Vendor & Third-Party AI**
Due diligence for AI-enabled solutions

06 **Roadmap**
30-day quick wins and 12-month plan

Optiri

The Three AI Risk Categories You Must Address

Shadow AI

Unauthorized use of consumer AI tools

- ChatGPT, Claude, Gemini
- Copilot in browsers
- AI writing assistants
- Image generators
- Personal accounts on work devices

Risk: Data leakage to training sets, no audit trail, compliance violations

Embedded AI

AI built into vendor solutions

- Loan origination/underwriting
- Fraud detection systems
- Security tools (SIEM, EDR)
- Member service chatbots
- Collections optimization
- Marketing/segmentation

Risk: Fair lending liability, model opacity, silent updates, vendor lock-in

Agentic AI

AI that takes autonomous action

- Auto-decisioning loans
- Automated remediation
- Multi-step workflows
- AI coding assistants
- Autonomous customer service
- Security orchestration

Risk: Cascading errors, no human review, accountability gaps, speed outpaces oversight

What is Shadow AI?

Shadow AI refers to the use of AI tools and services without IT department knowledge or oversight—bypassing security controls and governance.

Common Examples:

- ChatGPT, Claude, Gemini via personal accounts
- Browser-based AI extensions
- Mobile AI apps on personal devices
- Copying member data into AI prompts

RED FLAGS TO WATCH

- Unusual clipboard activity
- Traffic to AI domains
- AI-characteristic phrasing
- Unapproved browser extensions

⚠ **Key Insight:** Shadow AI can extract sensitive data with a single prompt.

The Shadow AI Problem

Your employees are using AI - whether you know it or not

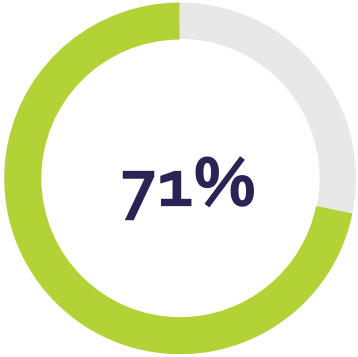

of office workers use AI **without IT approval**

Source: 2025 Research

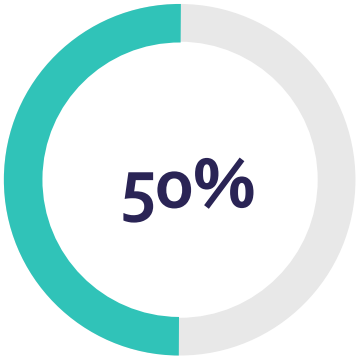

of employees are shadow AI users; **46% would refuse to stop** if banned

Source: Software AG 2024 (6,000 workers)

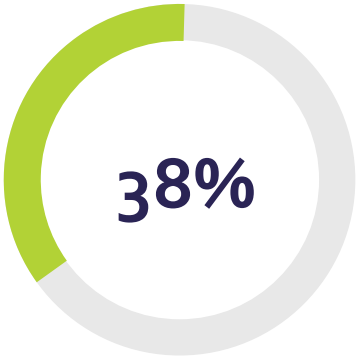

share **sensitive work info** with AI without employer knowledge

Source: CybSafe/NCA 2024 (7,000 respondents)

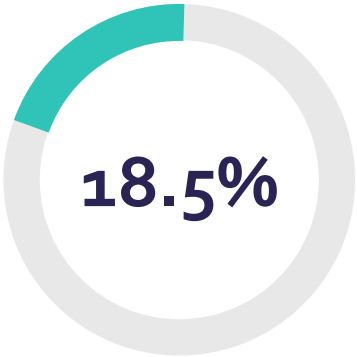

are **aware of their company's AI policy**

Source: 2025 Survey (12,000 workers)

Embedded AI: The AI You Already Have

The Liability Problem

When a vendor's AI makes a discriminatory lending decision, YOU own the fair lending liability.
"The vendor's model did it" is not a defense under ECOA or fair lending laws

Where AI is Already Embedded in Your Stack:

Lending & Credit
Underwriting, pricing, risk scoring, adverse action

Collections
Skip tracing, payment propensity, contact optimization

Member Service
Chatbots, IVR, knowledge search, next-best-action

Security Tools
SIEM, EDR, UEBA, threat intelligence, SOC automation

Fraud & BSA/AML
Transaction monitoring, behavioral analytics, SAR filing

Marketing
Segmentation, targeting, personalization, offers

Hidden Risks:

- **Silent model updates** – Vendor changes behavior without notice
- **Model drift** – Accuracy degrades over time
- **Black box** – Can't explain decisions to examiners
- **Training data** – May include biased historical data

Optiri

Agentic AI: The Emerging Frontier of Risk

What Makes AI "Agentic"?

Unlike traditional AI that provides recommendations, agentic AI takes autonomous action without human approval.

Examples:

- Auto-decisioning – Approves/denies loans without review
- Security orchestration – Isolates endpoints automatically
- Workflow automation – Multi-step processes across systems
- AI coding tools – Writes and deploys code

Required Controls:

- Human-in-the-loop for high-risk decisions
- Kill switches – instant halt capability
- Action limits – cap before human review

Unique Risks:

Cascading Errors

Chain reactions across systems

Accountability Gap

Who's responsible?

Scope Creep

Actions beyond boundaries

Prompt Injection

Manipulated into harm

Speed outpaces oversight. One bad decision becomes 500 before anyone notices.

Why Credit Unions Are High-Value Targets

300× more likely to be targeted vs. other industries
IBM Security

$5.9M Average breach cost in financial services

Contributing Factors:

- **Limited Security Resources**
  Smaller IT teams, often no dedicated security staff, competing priorities
- **Member Trust = High-Value Data**
  Complete financial profiles, SSNs, account access – premium on dark web
- **Heavy Vendor Dependence**
  Core processors, fintechs, and service providers all adding AI features
- **Productivity Pressure**
  Staff eager to use AI tools to do more with less – with or without approval

NCUA INCIDENT DATA (Sept 2023 – May 2024)

892 Cyber incidents

73% Third-party involved

Member Data at Risk:
SSNs, account numbers, credentials, loan applications, internal procedures

The Data Exposure Crisis

Corporate data flowing to AI tools at unprecedented rates

485%

increase in corporate data sent to AI tools

March 2023 - March 2024 | Cyberhaven Labs

46%

experienced internal data leaks through GenAI

Cisco 2025

22%

of files uploaded contain sensitive info

Netskope Q2 2025

4.7%

of financial firm employees put data in AI

March 2024

Regulatory Landscape for AI

NCUA Expectations

- Risk assessment for AI/ML adoption
- Board oversight and reporting
- Vendor due diligence for AI services

NIST AI RMF

- GOVERN – Establish accountability
- MAP – Identify AI uses
- MEASURE – Assess risks
- MANAGE – Mitigate and monitor

FFIEC (Emerging)

- Model risk management
- Third-party AI oversight

Fair Lending Focus

ECOA applies regardless of human or algorithm. Examiners ask:

- How are AI lending models validated?
- What testing for disparate impact?
- Can you explain adverse actions?

Key Takeaway

Regulators don't need new AI rules. Existing model risk and vendor guidance already applies.

Examiners ARE asking about AI.

NIST AI Risk Management Framework

GOVERN

Cultivate organizational culture and establish accountability

- Board oversight
- AI ethics principles
- Roles & responsibilities
- Policy framework

MAP

Identify and understand AI systems in your environment

- AI inventory (all 3 types)
- Use case documentation
- Data flow mapping
- Stakeholder identification

MEASURE

Assess, analyze, and track AI risks

- Risk assessments
- Bias testing
- Performance metrics
- Compliance validation

MANAGE

Mitigate, monitor, and respond to risks

- Controls implementation
- Incident response
- Continuous monitoring
- Ongoing improvement

Optiri

The Governance Gap

<⅓

have AI governance frameworks

1 in 5

advanced maturity

70% use GenAI tools

vs

15% implemented policies

Security Leaders Know It

69%

suspect employees using public GenAI

Building Your AI Governance Structure

Cross-Functional Team

- **IT/Security** – Controls
- **Compliance** – Regulatory
- **Risk** – Assessment
- **Legal** – Contracts
- **Business** – Use cases

Board Oversight

- Define risk appetite
- Quarterly AI reports
- Major use case approval
- Policy sign-off
- Escalation path

Essential Policies

- Acceptable Use Policy
- Data Classification
- Approved Tools List
- Vendor Criteria
- Incident Response

Step 1: Build Your Comprehensive AI Inventory

You can't govern what you don't know exists. Your inventory must include ALL three AI categories:

Shadow AI Discovery	Embedded AI Discover	Agentic AI Discovery
• Network traffic analysis • Browser extension audit • Employee surveys • CASB/DLP alerts	• Vendor questionnaires • Contract AI clause review • Product documentation • Vendor briefings	• Automation workflow review • RPA/orchestration platforms • Auto-decisioning systems • Security automation

Minimum Inventory Fields:

- System name and vendor
- AI category
- Business owner
- Data inputs/outputs
- Risk tier
- Approval status
- Last review date
- Human-in-the-loop req.

Technical Controls for AI Risk Mitigation

Detection

- DNS/Web filtering – Block or alert on AI domains
- DLP rules – Detect sensitive data to AI services
- CASB policies – Monitor cloud AI usage
- Endpoint agents – Track AI application usage
- Browser controls – Manage AI extensions

Prevention

- Enterprise AI tools – Approved alternatives
- Access controls – Role-based AI permissions
- Data classification – Prevent restricted data uploads
- Network segmentation – Isolate AI systems
- API gateways – Control AI API access

Agentic AI Controls

- Action limits – Cap transactions before human review
- Kill switches – Instant halt capability
- Audit logging – Every AI action recorded
- Anomaly detection – Alert on unusual patterns
- Rollback capability – Undo AI actions

Vendor AI Due Diligence

Contract Must-Haves

- **Model Change Notification** - 30-day advance notice of material AI model changes
- **Audit Rights** - Right to audit AI systems and request model documentation
- **Data Usage Limitations** - Prohibit use of your data to train models for others
- **Liability Allocation** - Clear terms for AI-related errors, fair lending violations
- **Data Deletion upon Termination** – Set deletion timeline and confirmation terms
- **Breach notification timelines**
- **Data Residency** – Limit data transmission, processing, and storage to within United States jurisdiction
- **SLA** – Establish performance metrics

Add These to Every Vendor Questionnaire:

1. Does your product use **AI, ML,** or **automated decision-making**?
2. What data is used to **train/operate** the AI models?
3. Is our data used to **train models** shared with **other customers**?
4. How are AI model updates **communicated** and **validated**?
5. Can you explain **model decisions** for **regulatory examination**?
6. What testing for **bias/disparate** impact has been **performed**?
7. Does the AI take **autonomous actions** or only **recommendations**?

30-Day Quick Start Plan

Actions you can take starting tomorrow

Week 1

- Survey staff on current AI use
- Review network logs for AI domains
- List all vendors – flag those with AI
- Brief executive team on AI risks

Week 2

- Draft acceptable use policy
- Identify oversight team members
- Create AI vendor questionnaire
- Review lending system AI features

Week 3

- Enable DLP alerts for AI services
- Evaluate enterprise AI options
- Hold first oversight team meeting
- Inventory agentic/auto-decisioning

Week 4

- Finalize and publish AI policy
- Send vendor AI questionnaires
- Brief board on findings
- Create 12-month roadmap

12-Month AI Governance Roadmap

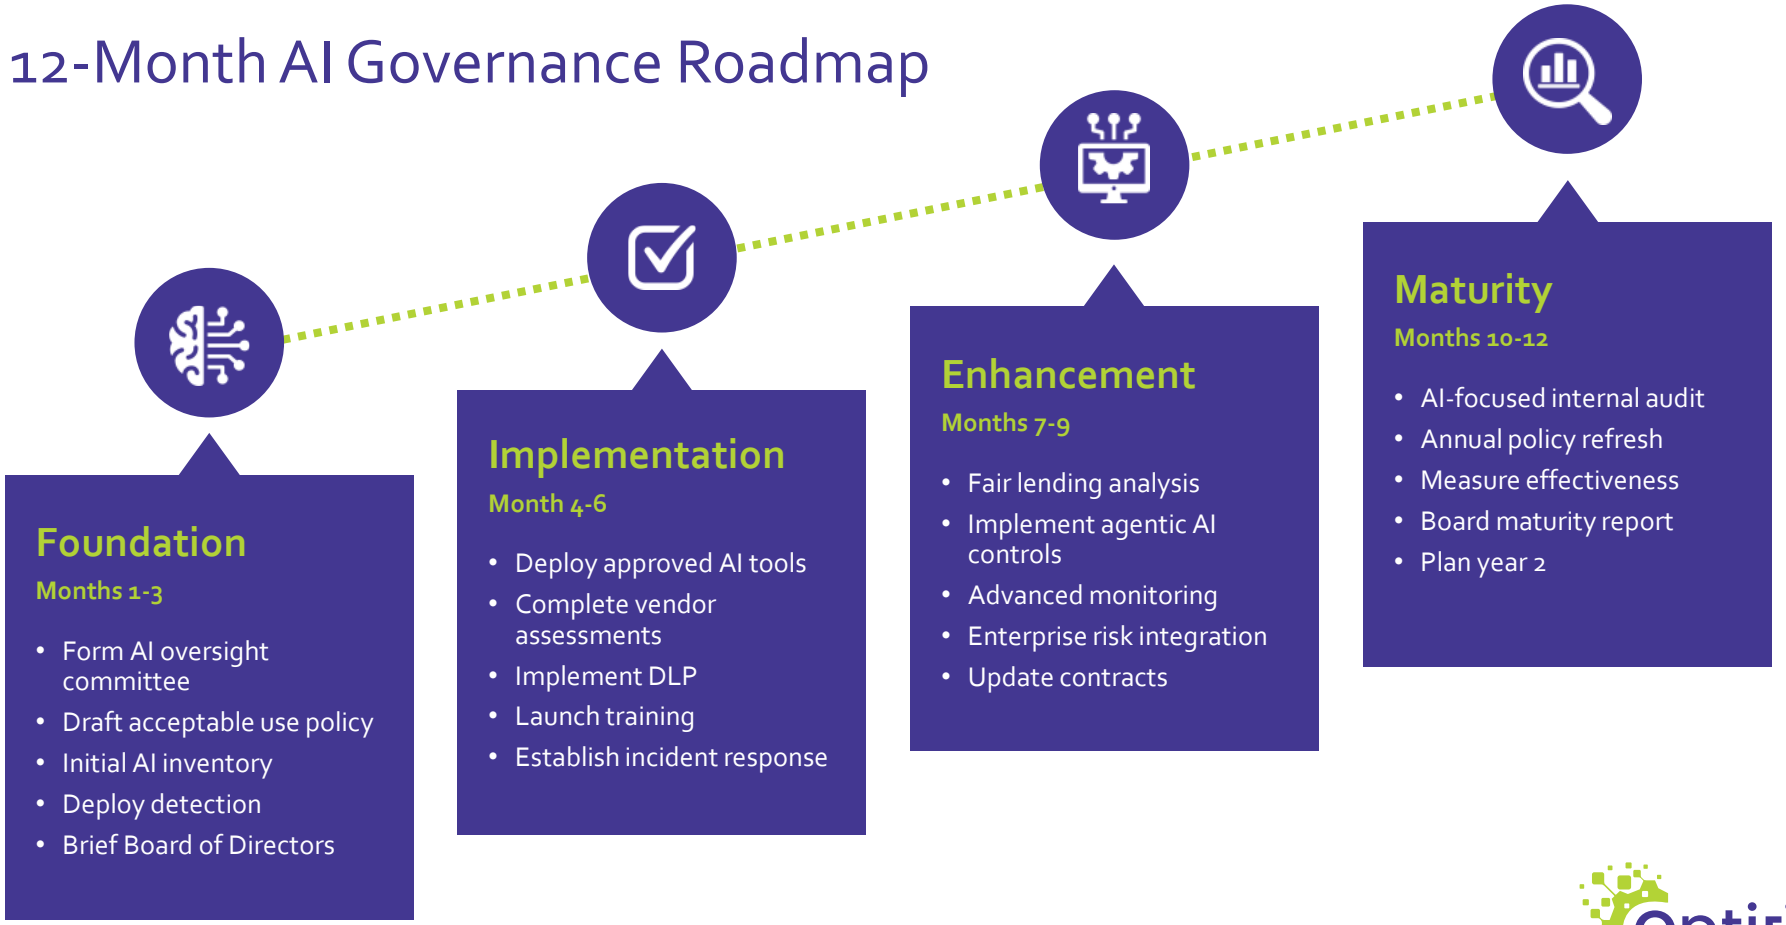

Key Takeaways

AI risk is broader than ChatGPT	Your inventory must include shadow AI, embedded vendor AI, and agentic AI systems
Vendor AI is YOUR liability	Fair lending, compliance, and member harm from vendor AI decisions are your responsibility
Agentic AI needs special controls	Human-in-the-loop, kill switches, and action limits are essential for autonomous systems
Regulators are already asking	NCUA examiners expect AI governance. Use NIST AI RMF as your framework
Start now – don't wait	Your 30-day plan starts tomorrow — you Use the 30-day plan. Prohibition isn't viable; governance is essential

Resources & Q&A

See handouts for additional resource directory

NCUA:	https://ncua.gov/ai
NIST:	https://www.nist.gov/itl/ai-risk-management-framework
CISA:	https://www.cisa.gov/ai

Thank you

Shane Butcher
sbutcher@optiri.com

Barry Lewis
blewis@optiri.com

AI Reference Guides

Optiri
A Trellance Company