

OFN42 · NEW ORLEANS · OCTOBER 22, 2026

Beyond Credit Risk

Building a Practical Enterprise Risk
Management Program for Emerging Risks

AI as an emerging risk. And a tool to manage risk.

SOC 2 Type 2 Certified

Governed AI: your content, your rules

OFN Ally

For questions or information reach out to suresh@salestable.ai

A free starter kit for community lenders

- Risk register starter
- Six questions for any AI vendor
- Sample one-page AI use policy
- Board and examiner readiness check
- Five staff pulse questions

Starter kit: risk register

RISK REGISTER STARTER (EXAMPLE ROWS)

Risk	Workflow	Owner	Control	Check	Last review
Confident wrong answers	Credit memo first draft	Head of Credit	Sources shown; people review	Monthly	Oct 2026
Data exposure	All AI tools	Head of IT	Approved tools; access by role	Quarterly	Oct 2026
Fabricated documents	Borrower financials	Head of Credit Ops	Verify at the source	Quarterly	Oct 2026
Unapproved actions	Covenant monitoring agent	Head of Portfolio	Agent flags only; a person acts	Monthly	Oct 2026
Unfair outcomes	Grant scoring	Head of Programs	Scheduled checks across applicant types	Quarterly	Oct 2026

One owner per row. Review each row on a fixed schedule.

Six questions to ask any AI vendor

THE QUESTION	WHY IT MATTERS	WHAT A GOOD ANSWER LOOKS LIKE
1. Where is our data stored, and who can see it?	You stay accountable for borrower and staff data, even inside a vendor's system.	A named region and cloud, encryption at rest and in transit, access by role, and access logs you can request.
2. Is our data ever used to train models, yours or anyone's?	Training can expose your data to others and cannot be undone.	A clear "no" in the contract, covering the underlying model provider too.
3. Can we see a current SOC 2 Type II report and your list of sub-processors?	It shows controls were tested over a period of time, and who else touches your data.	A report from the last 12 months, exceptions explained, and notice before sub-processors change.
4. How does each answer show its source, and how do we review and correct it?	Staff need to check answers fast, and examiners will ask how.	A link to your own document on every answer, and a simple way to flag and fix a wrong one.
5. Which actions can the AI take on its own, and where does a person approve?	Autonomy is where operational risk lives.	A written list of allowed actions, approval points you set, and a log of every step.
6. What happens to our data and our configuration if we leave?	Lock-in and data left behind are both third-party risks.	Export in a standard format, certified deletion within a set time, and help with the handover.

Sample: one-page AI use policy

[Organization] · Owner: [Role] · Approved by: [Board or committee], [Date] · Reviewed at least once a year

1. Scope

Applies to all staff, contractors and board members using AI for [Organization] work, including AI features inside software we already use.

2. Approved tools

Use only tools on the approved list: [list]. New tools, and new AI features in existing tools, need approval from [Role] before use.

3. Protect borrower and staff data

Never enter nonpublic personal information (Social Security or account numbers, credit reports, tax returns, financial statements) into a tool not approved for it. This supports our GLBA Safeguards Rule program.

4. A person stays responsible

AI output is a draft. The person who uses it owns its accuracy. A named person reviews AI-assisted credit memos, grant decisions, funder and board reports, and anything sent outside [Organization].

5. Credit decisions

AI does not approve, deny or price a loan on its own. Adverse action notices state the specific reasons for the decision, as Regulation B requires. Any AI use in credit decisions needs [Committee] review first, including fair lending testing.

6. Vendors

Before a vendor's AI touches our data: answer the six vendor questions, confirm in the contract how our data is used and whether it trains models, and add the tool to the risk register.

7. Payments and fraud

Verify every payment, wire or account change request through a known phone number or a second channel, even when the voice, video or document looks real.

8. Records

Note in the file when AI helped prepare a credit, grant or reporting decision, and keep the record under our retention schedule.

9. Report problems

Report wrong outputs, suspected data exposure or misuse to [Role] the same day.

10. Training

Everyone completes AI training at onboarding and every year. [Role] reviews this policy and the AI rows of the risk register at least once a year.

Sample for discussion, not legal advice. Review with counsel before adopting. Depository CDFIs should also align with their regulators' third-party risk guidance. Housing loans are also covered by the Fair Housing Act.

Five questions your board and examiners will ask

THE QUESTION

- 1 **Where are we using AI, and who owns each use?**
- 2 **Where do the AI's answers come from?**
- 3 **What can the AI do without a person?**
- 4 **What do our vendors do with our data?**
- 5 **How do we know outcomes are fair and inputs are real?**

A READY ANSWER

A risk register row for every AI use, with a named owner and a review date.

Our governed policies and documents. Every answer shows its source.

Draft and flag only. People approve credit, payments and anything sent outside.

Answered in writing: storage, training, SOC 2 report and exit terms.

Scheduled fairness checks, and source verification for borrower documents.

If you can answer these five, **your program is ready for the conversation.**

Starter kit: five staff pulse questions

THE QUESTION	WHY ASK IT	NEXT STEP IF THE ANSWER IS NO
1. Do you know which AI tools you may use?	Tells you how much shadow use to expect.	Publish a one-page list of approved tools this month.
2. Has your manager shown you a real use?	People adopt what their managers use.	Ask each manager to demo one approved use in a team meeting.
3. Do you know when a person must review AI output?	This is your main control. It only works if people know it.	Add review rules to the one-page policy and to onboarding.
4. Do you know who to ask when an answer looks wrong?	Wrong answers that go unreported never get fixed.	Name a contact and a simple way to flag a bad answer.
5. Has AI removed a task you disliked?	Shows whether AI is actually helping.	Find one repetitive task per team and pilot AI there.

Run it every quarter. Watch the trend over time.

AI READINESS & ENABLEMENT FOR CDFIs

Institutional knowledge, ready when your team needs it

SalesTable centralizes policies, procedures, and institutional knowledge into one AI-guided system, so every team, from lending to operations, can onboard faster and execute more consistently.

SOC 2 Type 2 Certified

Governed AI: your content, your rules

OFN Ally

For questions or information reach out to suresh@salestable.ai or visit <https://impact.salestable.ai/>

Suresh Madhuvarsu

Founder & CEO | suresh@salestable.ai

October 2026