Multimodal.

AI Vendor Due-Diligence Questionnaire for Credit Unions

A 40-question evaluation for any AI vendor touching member data, lending decisions, or core systems.

Why This Questionnaire Exists

Across 445 prospect conversations at 144 financial institutions, including 70 credit unions, Multimodal's 2026 Agentic AI Field Report found that adding an AI vendor functions as a regulatory event: every new vendor triggers a third-party risk review, a vendor management policy update, a board reporting line item, and in many cases an examiner question at the next audit. The same dataset shows auditability overtaking accuracy as the deciding evaluation criterion. This questionnaire turns that regulatory event into a documented, repeatable process.

How to Use This Questionnaire

1. **Classify the vendor first.** Take a risk-based approach. An AI tool drafting internal copy does not need the same scrutiny as a system touching member data, lending decisions, or core systems. Send the full questionnaire to high-risk vendors only; use Sections 1, 2, and 6 as a short form for low-risk tools.
2. **Require documented evidence.** Every answer should include a document: a SOC 2 report, a data flow diagram, a validation report, or a contract clause. Verbal assurances do not survive an exam.
3. **Score the answers, not the demo.** Use the same questionnaire for every vendor on your shortlist so your procurement team can make standardized, side-by-side comparisons.
4. **File everything.** The completed questionnaire belongs in your vendor-management file. NCUA cannot examine your technology vendors, so this documentation is the record your examiner reads.
5. **Repeat annually.** Vendor due diligence is a continuous obligation. Resend this questionnaire for re-attestation every year and after any material change to the vendor's models, sub-processors, or data practices.

Vendor Information

Field	Response
Vendor name	
Product evaluated	
Use case at our credit union	
Risk classification (High / Medium / Low)	
Date sent	
Date returned	
Reviewed by	
Decision / next step	

Section 1: Company Viability and Experience

What this verifies: *financial health, stability, and insurance. (NCUA 01-CU-20)*

#	Question	Vendor response	Evidence provided (Y/N)
1	*Provide audited financials or equivalent proof of financial health for the past two years.*		
2	*How many financial institution clients do you serve, and how many are credit unions? Provide two references.*		
3	*What insurance do you carry (cyber liability, technology E&O, commercial), and at what limits?*		
4	*Who owns the company, and what is your funding runway?*		
5	*Describe your disaster recovery and business continuity plans, including tested recovery time objectives.*		

What a strong answer looks like:
Documents, not adjectives. Audited statements, a certificate of insurance, named references, and a tested DR plan.

Section 2: Data Governance and Security

What this verifies: *where member data goes, and how it is protected. (NCUA AI guidance; GLBA safeguards)*

#	Question	Vendor response	Evidence provided (Y/N)
6	*Map the full data flow: where is our data received, processed, stored, and backed up, and in which jurisdictions?*		
7	*Is our data encrypted in transit and at rest? Specify standards.*		
8	*Is our member data ever used to train your models or any third party's models? If contractually barred, cite the clause.*		
9	*What are your data retention and deletion policies, including backups, and how is deletion verified?*		
10	*Describe your access controls: who at your company can see our data, under what approvals, with what logging?*		

What a strong answer looks like:
A data flow diagram, an encryption specification, a written retention schedule, and a contractual bar on training with customer data.

Section 3: Sub-Processors and Fourth Parties

What this verifies: *the vendor's own supply chain. (NCUA 07-CU-13 / SL 07-01)*

#	Question	Vendor response	Evidence provided (Y/N)
11	*Provide your current list of material sub-processors, including cloud and model providers.*		
12	*What notice do we receive before you add or change a material sub-processor?*		
13	*Do your data-protection obligations flow down to every sub-processor in writing?*		
14	*Which foundation models power your product, and what happens to our data under those providers' terms?*		
15	*If a sub-processor suffers a breach, what is your notification commitment to us?*		

What a strong answer looks like:
A named sub-processor list with written flow-down terms. Treat any claim of "zero sub-processors" with scrutiny; modern AI products rely on cloud and model providers.

Section 4: Model Governance and Explainability

What this verifies: *validation, bias testing, and transparency. (NIST AI RMF, Measure; GAO 2025)*

#	Question	Vendor response	Evidence provided (Y/N)
16	*How are outputs generated, and can each output be traced back to its inputs and reasoning?*		
17	*How do you validate model performance? Describe automated testing against benchmarks and test data, plus human review.*		
18	*What bias testing do you perform, what metrics do you use, and how often? Provide the latest results.*		
19	*What are the documented, known limitations of your models, and where do they fail?*		
20	*How do you evaluate new model versions before release? Describe regression checks and side-by-side comparisons.*		

What a strong answer looks like:
Validation reports, bias testing metrics, and documented limitations handed over without being pressed. Black-box answers are unacceptable where fair-lending laws apply.

Section 5: Agentic AI Controls

What this verifies: *autonomy limits, human oversight, and audit trails. (NIST AI RMF, Govern and Manage)*

#	Question	Vendor response	Evidence provided (Y/N)
21	*What actions can the system take autonomously, and what are the hard boundaries?*		
22	*Where are the human review checkpoints, and can we configure human oversight per workflow?*		
23	*Do audit logs capture every prompt, output, action, and access event? For how long, and can we export them?*		
24	*What tool, API, and system permissions does the agent hold, and how are they scoped?*		
25	*Is there a rollback and kill-switch mechanism, and who can trigger it?*		

What a strong answer looks like:
A live demonstration of the audit log and configurable human checkpoints. If the vendor cannot show you what its agent did and why, you cannot show your examiner either.

Section 6: Compliance and Certifications

What this verifies: *security certifications and regulatory posture. (NCUA AI guidance)*

#	Question	Vendor response	Evidence provided (Y/N)
26	*Provide your current SOC 2 Type II report or ISO 27001 certificate.*		
27	*How does your AI governance program align with the NIST AI RMF? Provide documented policies for responsible design and deployment.*		
28	*How do you track and comply with emerging AI regulations, including state AI laws, and GDPR and the EU AI Act where they apply?*		
29	*Will you support our exams: examiner questions, documentation requests, audit rights?*		
30	*Have you had a security incident or regulatory action in the past five years? Describe it and the remediation.*		

What a strong answer looks like:
The certification itself plus written AI governance policies with a named accountable owner. SOC 2 is the floor, not the finish line.

Section 7: Contract and Exit Terms

What this verifies: *AI clauses, SLAs, data return, and lock-in risk. (NCUA 07-CU-13; GAO 2025)*

#	Question	Vendor response	Evidence provided (Y/N)
31	*Is the uptime SLA written into the order form, with remedies? What figure do you commit to?*		
32	*What is your breach notification window? (72 hours after confirmation is the common documented standard.)*		
33	*Can you change AI-related terms unilaterally? What notice and objection rights do we have?*		
34	*What are the data return and export terms at exit: format, timeline, cost, and verified deletion afterward?*		
35	*What proprietary components would make switching vendors difficult, and how do you mitigate lock-in?*		

What a strong answer looks like:
Commitments in the contract, not in the sales deck. A clean exit strategy with clear data export terms, negotiated before signature, is the only leverage you keep afterward.

Section 8: Implementation and Ongoing Monitoring

What this verifies: *integration, reporting, and continuous oversight. (NCUA 2026 supervisory priorities)*

#	Question	Vendor response	Evidence provided (Y/N)
36	*How does your system integrate with our core and LOS, and which named integrations run in production today?*		
37	*Can we run a time-boxed pilot with defined success metrics before full deployment?*		
38	*What performance reporting do we receive in production, at what cadence, and against which metrics?*		
39	*How will you notify us of material changes to models, sub-processors, or data practices?*		
40	*Will you complete an annual re-attestation of this questionnaire?*		

What a strong answer looks like:
Named production integrations, a pilot against pre-agreed success metrics, and a written commitment to material-change notice and annual re-attestation.

Red-Flag Scorecard

Check any that apply. One flag starts a negotiation. A pattern of flags ends one.

- ☐ Refuses to name sub-processors or foundation model providers
- ☐ Makes absolute claims ("zero retention, period") with no schedule to back them
- ☐ Has no audit logs, or logs it cannot export for you
- ☐ Quotes SLAs and breach windows verbally but leaves them out of the contract
- ☐ Provides no bias testing results, model limitations, or validation evidence
- ☐ Resists a time-boxed pilot or defined success metrics

Attestation

Field	Response
Vendor representative (name, title)	
Signature and date	
Credit union reviewer (name, title)	
Review date	
Next re-attestation due	

Field data cited from Multimodal's 2026 Agentic AI Field Report: 10 Mid-Market Insights, based on 445 prospect conversations at 144 financial institutions, October 2025 to May 2026.

This questionnaire is provided by Multimodal for informational purposes and does not constitute legal advice. Consult your counsel and your vendor-management policy before relying on it for a specific engagement.

Multimodal builds AgentFlow, an agentic AI platform for credit union lending and document workflows. See how AgentFlow answers these 40 questions at multimodal.dev.