



Bank AI Deployment Map and Evidence Checklist

Sort your bank's AI use cases into three deployment zones, then score the evidence file an examiner will ask for in each one.

How to use this checklist

Work through each page with the system's model owner and a compliance lead, and score the file as it stands today. The scores give both teams one view of where the evidence is thin.

1. List your AI use cases in the inventory table and place each one in a zone, using the deployment map above it as the reference.
2. Answer the classification questions on that zone's page. When an answer points to another zone, move the use case there.
3. Score every artifact and vendor question on the 0 to 3 scale below.
4. Add up each zone's score and read the result on the summary page.

Scoring scale

- 0, missing: no record exists.
- 1, draft: a record exists without a named owner.
- 2, complete: the record has a named owner and covers the current version of the system.
- 3, maintained: the record is also reviewed on a schedule, and the last review date is on file.

The deployment map

Placement follows the system's intended purpose, regardless of which model runs inside it.
A human reviewer adds a control and leaves the tier unchanged.

Use case	Zone	What applies	When
Fraud detection and transaction monitoring	Deploy	Outside Annex III point 5(b). GDPR Article 22 and FATF Recommendation 15 still apply	In force now
Internal and back-office automation	Deploy	Minimal risk. Scope on what the tool decides	In force now
Engineering and developer tooling	Deploy	Article 4 literacy duty only. Exposure is code confidentiality and dependency provenance	In force now
Customer-facing chatbots and servicing	Deploy with controls	AI Act Article 50 disclosure	Applies since August 2, 2026
Personalization and next-best-offer	Deploy with controls	AI Act Article 5(1)(a) and 5(1)(b) prohibitions	In force since February 2, 2025
AML risk profiling and customer risk rating	Deploy with controls	BSA/AML model risk statement rescinded April 17, 2026	Framework now self-set
Credit scoring and creditworthiness	Hold	Annex III point 5(b) and the Article 27 fundamental rights impact assessment (FRIA). GDPR Article 22 per Case C-634/21 applies now	December 2, 2027
Insurance risk assessment and pricing	Hold	Annex III point 5(c). Same high-risk stack	December 2, 2027
Generative and agentic AI in decisioning	Hold	Outside SR 26-2 scope. Annex III point 5(b) still applies to credit use	Framework pending

Your inventory

Use case	System owner	Zone	Why it belongs in that zone

Zone 1: Deploy

No new regime attaches, so existing controls carry the use case.

Classification questions

Question	If yes
Does any output feed an assessment of an individual's creditworthiness?	Move it to Hold under Annex III point 5(b).
Does the system run AML/CFT checks?	Record that it assesses no creditworthiness. The fraud exception excludes AML/CFT.
Can it block an account with no human review?	GDPR Article 22 applies when the effect is legal or similarly significant.
Does it evaluate staff performance or allocate tasks by behavior?	Treat that use as Hold under Annex III point 4.
Does a developer tool send code or customer data outside the bank?	Document the data path before wider use.
Does the system talk directly with customers?	Move it to Deploy with controls under Article 50.

Artifacts an examiner asks for

Artifact	Score (0 to 3)
Inventory entry with a named owner, and the AI Act tier and materiality tier as separate fields	
Intended-purpose statement naming fraud detection as the main purpose	
Data lineage for every feature, from source to model	
Validation report with benchmark performance and tested boundaries	
Alert and case trail showing why each alert fired and who closed it	
GDPR Article 22 record, with the route to human review	
FATF Recommendation 15 risk assessment, dated before go-live	
Monitoring and drift thresholds, with alerts that reach an owner	
Approved tool list and data path for each developer tool	
Article 4 training records for the staff who run the system	

Vendor-governance questions

Question	Score (0 to 3)
Does your use match the contract's stated purpose? If not, Article 25 can make the bank the provider.	
Where does the vendor process your data and code, and who can access them?	
Does the vendor's documentation map onto your inventory entry and validation report?	
Does your third-party risk program cover the vendor like any other supplier?	
For EU entities: does your DORA register of information list the vendor's ICT services?	
Who owns each artifact after the contract ends?	

Readiness score: artifacts ____ of 30 · vendor questions ____ of 18 · zone total ____ of 48 ·
owner _____ · review date _____

Zone 2: Deploy with controls

A named obligation applies, and the bank can meet it now.

Classification questions

Question	If yes
Does the system interact directly with people?	Article 50 disclosure has applied since August 2, 2026.
Did the bank build the generative system itself?	Article 50(2) marking applies. Systems on the market before August 2, 2026 have until December 2, 2026.
Can the output resemble an investment recommendation?	MiFID II suitability attaches. Set and test an advice boundary.
Does personalization run inside credit or debt-collection flows?	Run the Article 5(1)(a) and 5(1)(b) review at scoping.
Does any nudge rely on a customer's financial distress?	Redesign it. Article 5 applies in full, whatever the risk tier.
Does a risk rating decide who opens or keeps an account?	GDPR Article 22 applies to that decision.
Does a risk rating feed an individual's credit assessment?	Move it to Hold under Annex III point 5(b).

Artifacts an examiner asks for

Artifact	Score (0 to 3)
Disclosure record with the interface text and the date it went live	
Machine-readable marking of output from generative systems the bank built	
Advice boundary specification, with test results	
Route to a person for disputes and complaints, with usage logs	
Prohibited-practice review at scoping, with a documented line between relevance and pressure	
Bank-written tiering logic for monitoring and risk-rating models, replacing the rescinded 2021 interagency statement	
Validation report and monitoring thresholds for each risk-rating model	
GDPR Article 22 record for account decisions, with the route to human review	
Article 4 training records for the staff who run the system	

Vendor-governance questions

Question	Score (0 to 3)
Which party marks generated output, and where is that recorded?	
Can you change the disclosure text and the route to a person without a vendor release?	
Which signals drive the personalization engine, and can you exclude signs of financial distress?	
Does the vendor document risk-rating features and tiering well enough to defend at examination?	
Where does conversation data go, and how long does the vendor keep it?	
Who owns each artifact after the contract ends?	

Readiness score: artifacts ____ of 27 · vendor questions ____ of 18 · zone total ____ of 45 ·
owner _____ · review date _____

Zone 3: Hold

The full high-risk stack applies, or no framework covers the use case yet. Hold the launch and start the file now.

Classification questions

Question	If yes
Does it evaluate an individual's creditworthiness or set their credit score?	High-risk under Annex III point 5(b) from December 2, 2027.
Does it price or assess risk for individuals in life or health insurance?	High-risk under point 5(c), on the same terms.
Does the system profile natural persons?	The Article 6(3) exit is closed to it.
Does a generative or agentic system shape an individual's credit assessment?	Point 5(b) applies. SR 26-2 excludes it, so document your own framework.
Has the bank substantially modified a licensed system or changed its purpose?	The bank takes on provider duties under Article 25.
Does a decision trigger a US adverse-action notice?	Regulation B requires specific principal reasons.

Artifacts an examiner asks for

Artifact	Score (0 to 3)
FRIA under Article 27, due before first use. It can cross-reference an existing data protection impact assessment.	
GDPR Article 22 safeguards, which apply now under Case C-634/21	
Per-decision explanation that yields the principal reasons for each adverse-action notice	
Bias testing against defined groups, with documented thresholds and a repeat schedule	
Data governance and lineage under Article 10, with bias checks before deployment	
Technical documentation under Article 11 and Annex IV	
Validation report on accuracy and robustness under Article 15	
Human oversight log showing override authority in use, under Articles 14 and 26	
Automatic logs under Article 12, retained under Article 26	
Monitoring and drift thresholds under Articles 72 and 26	
For generative decisioning: a substitute framework, such as the NIST AI Risk Management Framework or ISO/IEC 42001	

Vendor-governance questions

Question	Score (0 to 3)
Has your team checked the provider's technical documentation against Annex IV?	
Will the provider's conformity assessment and registration cover your purpose by December 2, 2027?	
Would your planned changes make the bank the provider under Article 25?	
Can the model produce specific reasons for each adverse-action notice?	
Does the vendor share its bias-testing results and describe the data behind them?	
Who owns each artifact after the contract ends?	

Readiness score: artifacts ____ of 33 · vendor questions ____ of 18 · zone total ____ of 51 ·
owner _____ · review date _____

Readiness summary

Zone	Artifacts	Vendor questions	Total	Maximum	Share of maximum
Deploy				48	%
Deploy with controls				45	%
Hold				51	%

Reading the result

- 0 to 25%: the file does not exist yet. Start with the inventory entry, because every other record refers back to it.
- 26 to 50%: records exist in parts. Name an owner for each draft, then complete the records other artifacts cite, such as data lineage and the validation report.
- 51 to 75%: most records exist. Put review dates on them and test the file against a mock examination request.
- 76 to 100%: the file supports the use case. Keep the review schedule and re-score after each change listed below.

For Hold use cases, the score tracks preparation. The launch decision belongs to counsel and the model risk function.

Re-score after each of these

- November 16, 2026: the comment period closes on the US third-party risk management proposal, ahead of final guidance.
- December 2, 2026: the Article 50(2) grace period ends for generative systems already on the market.
- By the end of 2026: the European Commission's final guidelines on high-risk classification are due.
- Date not set: the US agencies' request for information on AI model risk.
- December 2, 2027: Annex III obligations start for credit scoring and insurance pricing.
- August 2, 2028: Annex I obligations start for AI inside regulated products.

Map your AI estate with our engineers

Bring your AI backlog, and our engineers will sort it into the three zones. They also flag where your architecture changes your regulatory role, and you receive the evidence pack each use case needs. Schedule a readiness call at sumatosoft.com/services/ai-readiness-assessment.

This checklist supports engineering readiness and does not constitute legal advice. Edge-case classification calls for counsel.

Sources:

1. [Federal Reserve SR 26-2 \(April 17, 2026\)](#)
2. [OCC Bulletin 2026-13](#)
3. [FDIC FIL-15-2026](#)
4. [Regulation \(EU\) 2024/1689 \(AI Act\)](#)
5. [Regulation \(EU\) 2026/1744 \(Digital Omnibus on AI\)](#)
6. [Regulation \(EU\) 2016/679 \(GDPR\)](#)
7. [Regulation \(EU\) 2022/2554 \(DORA\)](#)
8. [Court of Justice of the EU, Case C-634/21](#)
9. [Regulation B, 12 CFR 1002.9](#)
10. [FATF Recommendations](#)
11. [OCC Bulletin 2026-46 \(proposed third-party risk management guidance\)](#).



Thank you for your time!

Any questions? Drop us a line!

Headquarters

One Boston Place, Suite 2602
Boston, MA 02108, United States

Other ways to get in touch

info@sumatosoft.com
sumatosoft.com