

SR 11-7 & OCC 2011-12 AI Compliance Checklist

A practical checklist for extending your model risk management program to enterprise AI, LLMs, and agentic systems. Use it to self-assess readiness before your next exam, internal audit, or board review.

How to use this checklist

Work through each section with the accountable owner. Mark each control as **In place**, **Partial**, or **Gap**. Any Partial or Gap items on tier 1 or tier 2 AI systems should be tracked in your remediation plan with a named owner and target date.

1. Model inventory & scope

Control	In place	Partial	Gap	Owner / notes
Every LLM, embedding model, RAG pipeline, and agent in production is entered in the model inventory.				
Third-party foundation models (OpenAI, Anthropic, Google, Bedrock, Azure OpenAI) are recorded with provider, version, and access pattern.				
Each AI use case is assigned a materiality tier (high / medium / low) using existing MRM tiering criteria.				
Shadow AI discovery has been run in the last 90 days (SaaS logs, expense reports, browser telemetry).				
A named business owner and technical owner are assigned to every tier 1 and tier 2 AI system.				

2. Development documentation & conceptual soundness

Control	In place	Partial	Gap	Owner / notes
A model card exists for every production AI system, covering purpose, inputs, outputs, and known limitations.				
Training data, fine-tuning data, and RAG corpora have documented provenance and licensing.				
Prompt templates and system prompts are version-controlled in a central registry.				
Evaluation methodology is documented, including datasets, metrics, and pass/fail thresholds.				
Design decisions (model choice, RAG vs fine-tune, guardrails) have a written rationale reviewable by validators.				

3. Independent validation (effective challenge)

Control	In place	Partial	Gap	Owner / notes
Validation is performed by a function organizationally independent of the developers.				
Domain-specific evaluation sets exist and are refreshed on a documented cadence.				
Adversarial testing covers jailbreaks, prompt injection, and data exfiltration attempts.				
Bias, fairness, and disparate-impact testing is executed for customer-facing decisions.				
Hallucination and grounding rates are measured against ground truth for RAG systems.				
Validation findings, limitations, and approved use boundaries are captured in a signed validation report.				

4. Ongoing monitoring

Control	In place	Partial	Gap	Owner / notes
Production telemetry captures quality, drift, refusal rate, latency, and cost per outcome.				

Hallucination and safety metrics are monitored with alert thresholds.				
Provider model changes (weights, routing, deprecations) are tracked and trigger revalidation.				
PII and sensitive-data leakage detection runs on model inputs and outputs.				
Business KPIs tied to the AI system are monitored alongside model metrics.				
Revalidation cadence is defined per tier (e.g., annually for tier 1, on material change otherwise).				

5. Governance, roles & controls

Control	In place	Partial	Gap	Owner / notes
The board or a designated committee has approved an AI risk appetite statement.				
Roles are defined across 1LOD (developers), 2LOD (MRM / risk), and 3LOD (internal audit).				
Approval gates by materiality tier are documented and enforced before production release.				
Human-in-the-loop requirements are defined for high-impact or customer-facing decisions.				
Agent action scopes, tool permissions, and rate limits are reviewed before deployment.				
An AI incident response playbook exists and has been exercised in the last 12 months.				

6. Third-party & vendor model risk

Control	In place	Partial	Gap	Owner / notes
Foundation-model providers have completed vendor risk review (SOC 2, data handling, sub-processors).				
Contracts cover model change notice, indemnification, and data-use restrictions.				
Exit and portability plans exist for each critical model provider.				
Provider evaluations are refreshed at contract renewal or on material provider change.				
Data residency, retention, and training opt-out settings are documented per provider.				

7. Policies, standards & training

Control	In place	Partial	Gap	Owner / notes
The model risk management policy explicitly covers generative AI, agents, and third-party models.				
An acceptable-use standard for generative AI is published and communicated to all staff.				
Mandatory AI risk training exists for developers, validators, and business owners.				
Policy and standards are refreshed at least annually to reflect regulatory and technology change.				
Regulatory change monitoring includes SR 11-7, OCC 2011-12, and emerging AI-specific guidance.				

8. Audit & regulatory readiness

Control	In place	Partial	Gap	Owner / notes
Internal audit has AI within its audit universe and a defined coverage cycle.				
Model inventory, validation reports, and monitoring evidence are stored in a system of record.				
Regulatory exam requests for AI can be fulfilled from the system of record without ad-hoc collection.				
Board reporting includes AI risk metrics on a defined cadence.				
A tabletop exercise for an AI incident or examiner request has been run in the last 12 months.				

Next steps

1. Consolidate Partial and Gap items into a single remediation plan with owners and dates.
2. Prioritize tier 1 AI systems and any customer-facing or credit-related use cases.
3. Benchmark your program with the free Agentic AI Scorecard at agenticaiscorecard.org/scorecard.
4. For an outside-in review, book a discovery call at agenticaiscorecard.org/contact.

This checklist is provided for informational purposes only and does not constitute legal, regulatory, or compliance advice. Consult your regulators, legal counsel, and internal risk functions for guidance specific to your institution. © Agentic AI Scorecard.