OCC SUPER INTELLIGENCE FRAMEWORK.

On November 4, 2026 the OCC gave its examiners a sixteen-component framework for reviewing how banks use Super Intelligence, the federal government's new name for AI. It is an examiner resource, not a rule: **control intensity follows what an SI system could do if it fails, not how sophisticated the model is.**

EXAMINER GUIDE, NOT A RULE | CONSEQUENCE SETS THE TIER | NO SEPARATE SI RISK | SIX PILLARS, 16 PARTS

01 THE MECHANISM

An Examiner Resource

OCC Bulletin 2026-54 helps examiners review SI in credit, payments, fraud, service, operations, risk and markets. It sets no supervisory expectations or enforceable standards, and banks need not assess themselves with it. **Examiners tailor work to each use case's materiality.** [01] [02]

SI Is the Law's AI

Executive Order 14434 has agencies say Super Intelligence where they said AI; SI means the AI defined in 15 USC 9401(3). Traditional and non-generative models stay under model risk guidance OCC 2026-13, **so this framework mainly reaches generative and agentic SI.** [04] [05] [06]

Consequence, Not Cleverness

Control intensity can rise with what an SI agent's failure or misuse could do. Consequence, transfer of agency and human oversight drive the tier, not technical sophistication. **A simple rules-based agent that can move money may outrank a model that only drafts.** [02] [01]

02 THE FLOW

How Examiners Scope It CLASSIFY

1. **Start from purpose:** a documented, testable statement of what each system or agent may do, who or what may invoke it, which data, systems and actions it may touch, and when it must stop for a human. Activity outside it is scope drift, a control failure even when the output looks reasonable. [02]
2. **Score the consequence:** decision consequence, reversibility, customer and financial impact, autonomy, data sensitivity, explainability and provenance combine into an illustrative composite tier, from Tier 1 Critical to Tier 4 Low. [02]
3. **Scale the controls:** higher tiers may warrant tighter limits on autonomy, stronger controls, more direct human oversight, more independent validation and sturdier resilience. Community banks are examined in proportion to their size, complexity and risk profile. [02] [01]

What They Look For PROVE

1. **Oversight that is real:** human-in-the-loop means approval before an action; human-on-the-loop means monitoring with the authority and practical ability to intervene. Reviewers need competence, authority and capacity; bulk approval and alert fatigue hollow the control out. [02]
2. **Containment for a misfire:** resilience covers an agent that stays available while acting wrongly: circuit breakers, rate limits, least-privilege permissions, kill switches, a tested fallback process, and a separate workstream to remediate affected customers. [02]
3. **Evidence, not assertion:** repeated and adversarial testing, red teaming and hallucination measurement; records that rebuild inputs, outputs, tool calls, approvals and overrides; version control over models, prompts and permissions; audit with access to the underlying evidence. [02]

03 CORE PILLARS & COVERAGE

PILLAR	WHAT IT COVERS
Purpose	The actions the SI system or agent is permitted to execute: defined purpose and permitted use.
Risk	The consequences if an output or action is not fit for purpose: risk classification and criticality.
Control	What stops unauthorized or unsafe actions: human oversight, model risk, data, privacy, compliance, cybersecurity.
Oversight	Who monitors the system and who answers for it: governance, human oversight, monitoring, reporting.
Resilience	Detecting, containing, recovering from and remediating failures while critical processes continue.
Assurance	Evidence that controls keep working: testing, documentation, third-party risk, internal audit.

04 RISK TIERS COMPARED

ILLUSTRATIVE TIER	HUMAN CHECKS EACH ACTION?	ILLUSTRATIVE PROFILE AND USE CASES
Tier 4: Low	Yes	No customer or financial data, no direct action, fully human-controlled: productivity tools, code assistance, employee Q&A.
Tier 3: Moderate	Yes	Decision support with mandatory human-in-the-loop, limited data sensitivity, reversible: credit-memo summaries, research.
Tier 2: High	No	Direct action or high-weight support, human-on-the-loop: SI-assisted underwriting, agent-sent customer messages, AML alert disposition.
Tier 1: Critical	No	Direct action, little or no reversibility, high impact, autonomous: payment agents, real-time credit decisions, fraud-hold release.

OCC SI SUPERVISORY FRAMEWORK.

CLASSIFY · CONTROL · PROVE

05 PROGRAM OPERATIONS

Keep the inventory examiners will ask for. For each use case record the owner, provider, purpose, permissions, autonomy, deployment status, monitoring and fail-over; update it on every material change and reconcile it to incident records, validation results, audit findings and board reporting. [02]

Tier before you deploy, re-tier on change. Score all seven factors and watch aggregate exposure where use cases share a provider, data, permissions or infrastructure. A vendor adding SI features or swapping the model under a product you already run counts as a change. [02]

Govern agents like privileged users. Give each agent its own identity, least-privilege credentials, logged tool calls and rate limits; treat prompt injection, data and model poisoning, model extraction and excessive agency as part of the attack surface. [02]

Write SI into contracts and incident playbooks. Define SI incidents (hallucination, misfire, drift, manipulation) with internal escalation and notification thresholds, settle liability for model errors and notice of model changes with providers, and keep existing notice rules in view. [02] [07] [08]

06 CONTROL MAPPING

OCC 2026-13 MODEL RISK

The April 17, 2026 interagency guidance covers traditional and non-generative, non-agentic models and leaves generative and agentic AI out of scope. The SI framework carries model risk ideas to them: more frequent re-validation under drift, repeated sampling for non-determinism, and alternatives to champion-challenger. [06] [02]

THIRD PARTIES (2023-17)

The interagency third-party guidance still governs vendors. The framework adds identity and permission limits for provider agents in bank systems, SI incident definitions, contract terms for model-error liability, and notice before a vendor changes models or adds SI. [07] [02]

INCIDENT NOTICE (12 CFR 53)

Banks must notify the OCC no later than 36 hours after determining a notification incident occurred. The framework adds that banks may set internal SI thresholds by severity, duration, customer impact, financial exposure, legal requirement and systemic effect. [08] [02]

AI LEXICON (AIEOG)

The framework's lexicon reuses definitions from the AI Executive Oversight Group lexicon Treasury released on February 19, 2026, with "AI" changed to "SI": agentic SI, hallucination, prompt injection, human-in-the-loop, drift. [09] [02]

07 WHAT THIS DOES NOT SOLVE

Not a Rulebook

It sets no expectations or enforceable standards. **Existing law and safety and soundness still bind.** [01] [03]

Not a New Risk Stripe

SI is not its own risk category. **It shifts risk in existing ones, such as operational and credit.** [01] [02]

Not a Vendor Shield

Providers may build and host the model; **accountability stays with the bank.** [02]

Not Finished

The OCC calls it a living document. **Expect regular revisions as SI changes.** [02] [03]

08 MYTHS VS. REALITY

MYTH	REALITY
"The OCC just imposed new AI rules on banks."	It is an examiner resource with **no new requirements or enforceable standards**; banks need not self-assess with it. [01]
"Super Intelligence is a new class of technology."	Under Executive Order 14434 it is **the technology the law already calls AI**, renamed. [04] [05]
"The most advanced model is always the riskiest."	The tier follows consequence, agency and oversight; **a simple agent that moves money can be Tier 1**. [02]
"A human click on every output is enough oversight."	Examiners look for **competence, authority, capacity and meaningful review**; bulk approval breeds alert fatigue. [02]
"Buying the model moves the risk to the vendor."	**Accountability generally remains with the bank**, including for provider agents inside its systems. [02]

09 SPEAK THE PROGRAM — KEY TERMS

SUPER INTELLIGENCE (SI)

The federal term, since EO 14434, for **what 15 USC 9401(3) defines as AI**.

AGENTIC SI

SI that **decides and acts** in its environment without direct human intervention.

HUMAN-ON-THE-LOOP

A qualified human **monitors and can override or halt** a system already operating.

SCOPE DRIFT

Activity outside an approved purpose; **a control failure even if the result looks right**.

MISFIRE

An SI agent that **stays available while acting incorrectly**.

COMPOSITE TIER

Seven risk factors scored into **Tier 1 Critical to Tier 4 Low**.