WHITE PAPER · BANKING EDITION

HELIOS

From Expectation to Evidence

Comparing the OCC's Super Intelligence Supervisory Framework to the IIA's Artificial Intelligence Auditing Framework, and what it means for banks



Katie Granville
Senior Partner, Helios

CONTENTS

In this paper

Introduction	03
Key Takeaways	05
Two Frameworks, Complementary Purposes	06
Scope, Methodology, and Alignment at a Glance	07
Cornerstone One: Control	08
Cornerstone Two: Oversight	11
Cornerstone Three: Resilience	13
Cornerstone Four: Assurance	14
Closing the Gap: Building the Translation Layer	16
Questions to Prepare For · How Helios Can Help	17

The OCC Framework defines what should be controlled; the IIA Framework supplies how to test it. Neither was designed with the other in mind.

KATIE GRANVILLE, SENIOR PARTNER

INTRODUCTION

A new supervisory perspective on AI in banking

On November 4, 2026, the Office of the Comptroller of the Currency ("OCC") issued OCC Bulletin 2026–54, *Super Intelligence: Supervisory Framework* (the "OCC Framework"). The OCC Framework is a resource designed to assist OCC examiners in evaluating banks' use of "super intelligence" ("SI"), the term the OCC now uses for technologies previously described as artificial intelligence ("AI"), including generative and agentic systems deployed in normal-course business activities. It organizes examiner considerations into 16 components grouped within six pillars.

Purpose	Risk	Control	Oversight	Resilience	Assurance
Outside scope	Outside scope	In scope	In scope	In scope	In scope

The six pillars of the OCC Framework, and the four analyzed in this white paper.

The OCC Framework is expressly an examiner resource rather than a source of new supervisory expectations, enforceable standards, or prescriptive requirements. Nevertheless, it previews how examiners are likely to evaluate a bank's use of SI under existing standards, and it emphasizes that control intensity should increase with the potential consequences of an SI agent's failure or misuse.

A number of frameworks provide perspective on the expectations for assuring that AI structures are effectively designed and functioning as intended. These include COSO's *Achieving Effective Internal Control over Generative AI* (February 2026); NIST's AI Risk Management Framework (January 2023) and its Generative AI Profile (July 2024); ISO/IEC 42001, AI Management System (December 2023), ISO/IEC 23894, AI Risk Management Guidance (February 2023), and ISO/IEC 42005, AI System Impact Assessment (May 2025); and the OECD Recommendation of the Council on Artificial Intelligence, together with a growing body of global and state laws and regulations. COSO's guidance is among the most recent. COSO retained its framework but changed what it assumes: "validate the system once" becomes "validate the output continuously," and "control the application" becomes "control the capability, based on how much you rely on it." That shift echoes the OCC Framework's consequence-based scaling. The IIA Framework is another such resource.

INTRODUCTION (CONTINUED)

An expanding landscape of AI assurance frameworks

Jan 2023	**NIST** · AI Risk Management Framework
Feb 2023	**ISO/IEC 23894** · AI Risk Management Guidance
Dec 2023	**ISO/IEC 42001** · AI Management System
Jul 2024	**NIST** · Generative AI Profile
Sep 2024	**IIA** · Artificial Intelligence Auditing Framework (updated)
May 2025	**ISO/IEC 42005** · AI System Impact Assessment
Feb 2026	**COSO** · Achieving Effective Internal Control over Generative AI
Nov 2026	**OCC** · Super Intelligence: Supervisory Framework

Selected frameworks by publication date. The two frameworks compared in this paper are highlighted. The OECD Recommendation and the growing body of global and state laws and regulations are not shown.

The Institute of Internal Auditors' ("IIA") *Artificial Intelligence Auditing Framework* (updated September 2024) (the "IIA Framework") addresses a different question. Built on the IIA's Three Lines Model, it guides internal auditors on how to test AI controls. Because the Three Lines Model, spanning governance, management, and internal audit, is industry-agnostic and concerns the soundness of control, the IIA Framework is purposely broad and has been supplemented with training, papers, and deep dives on specific topics to remain current in the AI space. The two frameworks were never designed to operate together. Read side by side, however, they illuminate how internal audit can help institutions prepare for examination, and where gaps exist that may need to be addressed.

This white paper maps four of the OCC Framework's six pillars (Control, Oversight, Resilience, and Assurance) against the IIA Framework and considers the practical lessons the comparison holds for banks preparing for examination.

KEY TAKEAWAYS

Six points for boards, management, and audit committees

01 **Expectations versus methodology**

The OCC Framework is comprehensive as to *what* should be controlled but deliberately limited as to *how* those controls should be tested. The IIA Framework supplies a testing methodology, but not one designed with the OCC Framework in mind.

02 **A translation layer**

Closing that gap requires a deliberate "translation layer" that converts the OCC's control expectations into executable audit procedures and defined evidence.

03 **Principal gaps**

The most significant gaps (operational resilience, the quality of human oversight, privacy risks such as model memorization and inference, and hallucination-rate testing) are not addressed by ready-made techniques in either framework. Auditors may draw on the NIST AI Risk Management Framework or develop bespoke procedures.

04 **Areas of strongest alignment**

Alignment is strongest in governance and accountability, documentation and auditability, internal audit independence, and cybersecurity, where existing IIA techniques transfer directly or require only targeted extension.

05 **Enduring audit principles**

Whatever approach an institution adopts, the IIA's core principles of traceability, independence, and governance remain the foundation.

06 **Examination readiness**

Although the OCC Framework creates no new requirements, banks should expect examiners to use it as a reference point and should be prepared to evidence how each in-scope SI use case is controlled, overseen, made resilient, and assured.

DETAILED ANALYSIS

Two frameworks, complementary purposes

The OCC Framework tells examiners what controls should exist. The IIA Framework tells internal auditors how to test them. The distinction matters: an institution can satisfy one framework while remaining exposed under the other. A bank whose controls track the OCC's expectations but whose internal audit function cannot evidence their operation will struggle to demonstrate effectiveness at examination. Conversely, an internal audit plan built solely on the IIA Framework may test controls diligently yet overlook risks the OCC has specifically identified.

DEFINES WHAT TO CONTROL

OCC Super Intelligence Supervisory Framework

Issued	November 4, 2026
Reference	OCC Bulletin 2026–54
Audience	OCC bank examiners
Structure	16 components, six pillars

DEFINES HOW TO TEST

IIA Artificial Intelligence Auditing Framework

Updated	September 2024
Publisher	Institute of Internal Auditors
Audience	Internal auditors
Foundation	Three Lines Model

Four questions across the AI lifecycle

The four OCC pillars analyzed in this white paper, referred to here as cornerstones, correspond to the lifecycle questions that boards, management, and examiners are likely to ask of any SI deployment.

Control

PREVENTIVE SAFEGUARDS

Are effective safeguards in place to prevent unintended or inappropriate AI outcomes?

Oversight

ACCOUNTABLE OVERSIGHT

Do accountable individuals have the visibility, authority, and capability to intervene?

Resilience

CONTAINMENT AND RECOVERY

Can the organization contain, recover from, and remediate an AI failure?

Assurance

EVIDENCE AND VALIDATION

Can the organization demonstrate that its controls operate as represented?

Scope and methodology

The analysis covers the 14 OCC components that fall within the Control, Oversight, Resilience, and Assurance pillars; the Purpose and Risk pillars are outside its scope. For each component, the pages that follow set out (i) the OCC's control expectation; (ii) similarities in the IIA Framework; (iii) material differences between the two; (iv) the audit technique and the evidence production expectations the IIA Framework sets; and (v) Helios' commentary on the degree of alignment and the work required to close any gap.[1] Component numbering follows the OCC Framework.

Alignment at a glance

PREVENTIVE SAFEGUARDS Control	ACCOUNTABLE OVERSIGHT Oversight	CONTAINMENT AND RECOVERY Resilience	EVIDENCE AND VALIDATION Assurance
2.3 Human-in-the-Loop / Human-on-the-Loop Gap	2.4 Governance & Accountability Strong alignment	2.10 Operational Resilience Gap	2.9 Third-Party & Vendor Risk Partial alignment
2.5 Model Risk Management Extensions Partial alignment	2.12 Performance Monitoring Partial alignment		2.13 Testing Gap
2.6 Data Governance Partial alignment	2.16 Reporting & Transparency Partial alignment		2.14 Documentation & Auditability Strong alignment
2.7 Privacy Gap			2.15 Internal Audit & Independent Challenge Strong alignment
2.8 Compliance Gap			
2.11 Cybersecurity Strong alignment			

Strong alignment IIA technique transfers directly or with minimal adaptation

Partial alignment Sound foundation; targeted extension required

Gap No established technique; bespoke procedures required

Ratings summarize the Helios commentary for each component.

1 This comparison included only the IIA's AI Auditing Framework. The IIA has supplemental materials and training that provide deep dives into additional topics not included in this paper.

CORNERSTONE ONE · PREVENTIVE SAFEGUARDS

Control

Are effective safeguards in place to prevent unintended or inappropriate AI outcomes?

The Control pillar is the most extensive of the four pillars analyzed. It comprises six components: human-in-the-loop and human-on-the-loop oversight (2.3), model risk management extensions (2.5), data governance (2.6), privacy (2.7), compliance (2.8), and cybersecurity (2.11).

ALIGNMENT **1** Strong alignment **2** Partial alignment **3** Gap

2.3 Human-in-the-Loop / Human-on-the-Loop

Gap

OCC EXPECTATION
Competent reviewers; defined review triggers; span-of-control limits; override logs; reviewer independence.

IIA SIMILARITIES
The IIA requires role-based access and segregation of duties.

DIFFERENCES
The IIA has no dedicated human-oversight concept; oversight is folded into generic access controls.

IIA TECHNIQUE / EVIDENCE
Access reviews and segregation-of-duties testing. No technique is provided for assessing review quality or span of control.

HELIOS PERSPECTIVE A substantive gap; institutions must develop bespoke procedures for testing the quality of human review.

2.5 Model Risk Management Extensions

Partial alignment

OCC EXPECTATION
Model risk management adapted for drift; explainability tools; repeated and distributional testing; comparative and benchmark testing.

IIA SIMILARITIES
The IIA requires pre- and post-launch testing and refers to the NIST AI Risk Management Framework.

DIFFERENCES
The OCC names specific methods; the IIA remains generic and defers to NIST.

IIA TECHNIQUE / EVIDENCE
Confirmation of a validation plan, its execution, and a revalidation schedule.

HELIOS PERSPECTIVE Aligned at a high level; technical depth is drawn from NIST rather than the IIA.

CORNERSTONE ONE · CONTROL (CONTINUED)

2.6 Data Governance

Partial alignment

OCC EXPECTATION	IIA SIMILARITIES	DIFFERENCES	IIA TECHNIQUE / EVIDENCE
Data lineage; drift monitoring; synthetic-data review; access controls; evolving data architecture.	The IIA addresses data integrity, lineage, access, segregation, and review of vendor SOC reports.	The IIA addresses data governance generically, with no synthetic-data or drift-specific testing.	Lineage documentation, access logs, SOC reports, and data-quality tests.

HELIOS PERSPECTIVE A sound foundation; synthetic-data and drift testing must be added.

2.7 Privacy

Gap

OCC EXPECTATION	IIA SIMILARITIES	DIFFERENCES	IIA TECHNIQUE / EVIDENCE
Controls against model memorization and inference; retention limits; vendor data-use limits; cross-border controls.	The IIA requires data to be secure, private, and confidential, and vendor relationships to be reviewed.	The OCC identifies memorization and inference risk; the IIA has no corresponding concept.	Privacy policy and vendor contract review. No technique is provided for memorization or inference testing.

HELIOS PERSPECTIVE The most significant gap in this cornerstone, as an audit methodology cannot readily test a risk it does not identify.

2.8 Compliance

Gap

OCC EXPECTATION	IIA SIMILARITIES	DIFFERENCES	IIA TECHNIQUE / EVIDENCE
BSA/sanctions and consumer compliance controls; hallucination controls on automated representations.	The IIA requires legal and regulatory monitoring and compliance disclosures in reporting.	The OCC is banking-specific and names hallucination risk; the IIA is generic and cross-industry.	Compliance sign-offs and monitoring logs. No hallucination-specific test is provided.

HELIOS PERSPECTIVE An expected gap, given the IIA's cross-sector scope.

CORNERSTONE ONE · CONTROL (CONTINUED)

2.11 Cybersecurity

Strong alignment

OCC EXPECTATION

Identity and access management (including non-human agents); anti-poisoning controls; monitoring for prompt injection and model extraction.

IIA SIMILARITIES

The IIA addresses encryption, intrusion prevention and detection, prompt and response logging, penetration testing, and phishing training.

DIFFERENCES

The OCC targets SI-specific attacks; the IIA's controls reflect conventional IT security.

IIA TECHNIQUE / EVIDENCE

Penetration test reports, security logs, and training records.

HELIOS PERSPECTIVE Strong alignment; penetration testing should be extended to prompt injection and model extraction.

OBSERVATIONS · CONTROL

Alignment is strongest in cybersecurity, where the IIA's conventional IT security controls provide a sound base that can be extended to SI-specific attack vectors. Model risk management aligns at a high level, although technical depth must be sourced from NIST rather than the IIA. The widest gap is privacy: the OCC identifies memorization and inference risk, a concept the IIA Framework does not address. Human oversight presents a comparable gap, requiring institutions to design their own procedures for testing review quality and span of control.

CORNERSTONE TWO · ACCOUNTABLE OVERSIGHT

Oversight

Do accountable individuals have the visibility, authority, and capability to intervene?

The Oversight pillar addresses whether accountable individuals can observe SI behavior and intervene meaningfully. It comprises three components: governance and accountability (2.4), performance monitoring (2.12), and reporting and transparency (2.16).

ALIGNMENT **1** Strong alignment **2** Partial alignment

2.4 Governance & Accountability

Strong alignment

OCC EXPECTATION	IIA SIMILARITIES	DIFFERENCES	IIA TECHNIQUE / EVIDENCE
Board risk appetite; lifecycle approval gates; reviewer fluency; retained accountability; reporting.	The IIA's Governance domain (strategy, board involvement, and an AI leadership team) maps closely.	The OCC is gate-specific; the IIA is broader and includes internal audit as a governance participant.	Board minutes, AI leadership team charter, and governance policies.

HELIOS PERSPECTIVE The strongest point of alignment; auditors should nonetheless map the IIA domain to the OCC's specific lifecycle gates.

2.12 Performance Monitoring

Partial alignment

OCC EXPECTATION	IIA SIMILARITIES	DIFFERENCES	IIA TECHNIQUE / EVIDENCE
Calibrated metrics; rare-event observability; drift monitoring; monitoring of vendor changes.	The IIA requires defined key performance indicators (KPIs) and periodic reporting to leadership.	The OCC is technically granular (drift, rare events); the IIA remains at the level of generic KPI reporting.	KPI definitions and reporting packages. No test is provided for drift or rare-event observability.

HELIOS PERSPECTIVE Aligned at the reporting level only; testing requires technical, engineering-level expertise.

CORNERSTONE TWO · OVERSIGHT (CONTINUED)

2.16 Reporting & Transparency

Partial alignment

OCC EXPECTATION	IIA SIMILARITIES	DIFFERENCES	IIA TECHNIQUE / EVIDENCE
Escalation thresholds; material-change reporting; reconciliation of reports to the AI inventory, incidents, and audit findings.	The IIA requires bias, compliance, and incident disclosures in board reporting.	The OCC adds a reconciliation requirement that the IIA does not specifically address, but that any sound auditor would naturally employ.	Reporting parameters and sample reports evidencing required disclosures.

HELIOS PERSPECTIVE Substantive alignment on content; a reconciliation test should be added.

OBSERVATIONS · OVERSIGHT

Governance and accountability represents the strongest point of alignment in the entire comparison, as the IIA's Governance domain maps closely to the OCC's expectations. Performance monitoring, by contrast, aligns only at the reporting level; the OCC's focus on drift and rare-event observability calls for technically specialized testing that the IIA Framework does not supply. Reporting content aligns well, but auditors should add a reconciliation test linking management and board reporting to the AI inventory, incident records, and audit findings. In practice, experienced auditors would typically extend these concepts into additional testing methodologies that incorporate the elements described above.

CORNERSTONE THREE · CONTAINMENT AND RECOVERY

Resilience

Can the organization contain, recover from, and remediate an AI failure?

The Resilience pillar addresses how an institution contains, recovers from, and remediates an SI failure that occurs notwithstanding preventive controls. It comprises a single component: operational resilience (2.10).

ALIGNMENT **1** Gap

2.10 Operational Resilience

Gap

OCC EXPECTATION

Misfire containment; failover testing; recovery review; customer remediation; agent-interaction controls.

IIA SIMILARITIES

The IIA's general risk-treatment language addresses the concern only in general terms.

DIFFERENCES

The IIA has no dedicated component; resilience, emergency shutdown mechanisms, and remediation are not addressed.

IIA TECHNIQUE / EVIDENCE

No specific IIA technique; auditors rely on generic business continuity and disaster recovery practices, for which other guidance exists.

HELIOS PERSPECTIVE The largest gap in the comparison.

OBSERVATIONS · RESILIENCE

Resilience represents the largest gap in the comparison. The IIA Framework contains no dedicated resilience component, and emergency shutdown mechanisms ("kill switches"), the containment of erroneous actions, and customer remediation are not addressed. Auditors must therefore rely on generic business continuity and disaster recovery practices, supplemented by bespoke procedures tailored to SI failure modes and agent interactions.

CORNERSTONE FOUR · EVIDENCE AND VALIDATION

Assurance

Can the organization demonstrate that its controls operate as represented?

The Assurance pillar addresses how an institution demonstrates, to its board and to examiners, that its controls operate as represented. It comprises four components: third-party and vendor risk (2.9), testing (2.13), documentation and auditability (2.14), and internal audit and independent challenge (2.15).

ALIGNMENT **2** Strong alignment **1** Partial alignment **1** Gap

2.9 Third-Party & Vendor Risk

Partial alignment

OCC EXPECTATION	IIA SIMILARITIES	DIFFERENCES	IIA TECHNIQUE / EVIDENCE
Controls over the scope of provider agents; incident-notification terms; liability clarity; retesting upon change or exit.	The IIA requires SOC report review, right-to-audit provisions, service-level agreements, and defined third-party roles.	The OCC is agentic-specific; the IIA reflects traditional vendor risk management.	SOC reports, SLA reviews, and a third-party inventory.

HELIOS PERSPECTIVE A sound foundation; extend testing to agent permissions and change notification.

2.13 Testing

Gap

OCC EXPECTATION	IIA SIMILARITIES	DIFFERENCES	IIA TECHNIQUE / EVIDENCE
Full lifecycle testing; adversarial testing; red teaming; LLM-as-judge evaluation; hallucination testing.	The IIA requires pre- and post-launch testing and bias checks.	The OCC names advanced methods; the IIA extends only to confirming that testing was performed.	Test plans and results, and bias documentation. No red-teaming or hallucination technique is provided.

HELIOS PERSPECTIVE A key evidentiary gap for higher-tier use cases.

CORNERSTONE FOUR · ASSURANCE (CONTINUED)

2.14 Documentation & Auditability

Strong alignment

OCC EXPECTATION	IIA SIMILARITIES	DIFFERENCES	IIA TECHNIQUE / EVIDENCE
Decision-traceability records; version control; retention; documentation proportionate to risk tier.	The IIA identifies auditability as a core attribute of AI projects and requires traceable logs.	The OCC is more prescriptive as to content and tiering.	Tracing sampled outputs through logs back to source.

HELIOS PERSPECTIVE The best-aligned component; the IIA's technique transfers directly.

2.15 Internal Audit & Independent Challenge

Strong alignment

OCC EXPECTATION	IIA SIMILARITIES	DIFFERENCES	IIA TECHNIQUE / EVIDENCE
Independence controls; full evidentiary access; risk-based coverage; testing of control agents; board reporting.	This is the IIA's core subject: independence, third-line assurance, and the IIA Standards.	The IIA is generic; the OCC's concept of auditing the control agent itself is specific to agentic systems.	Audit charter, organizational chart, risk-based audit plan, and access records.

HELIOS PERSPECTIVE The strongest structural alignment; extend coverage to auditing AI-based controls themselves.

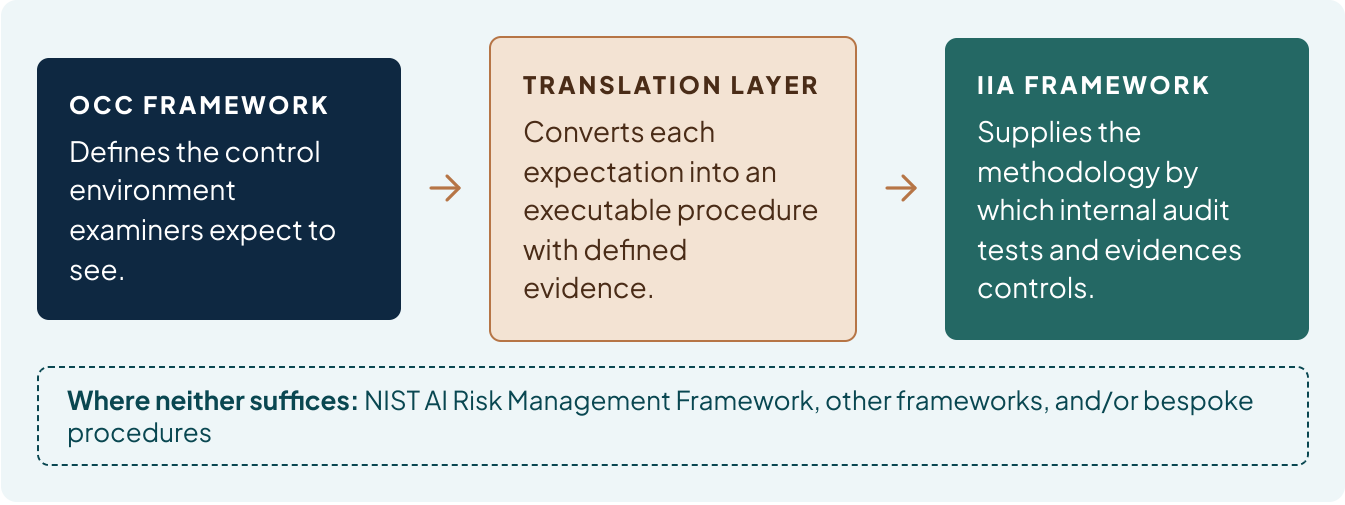
OBSERVATIONS · ASSURANCE

Documentation and auditability is the best-aligned component in the comparison, as the IIA's traceability technique transfers directly. Internal audit and independent challenge reflects the strongest structural alignment, although coverage must extend to AI-based controls themselves, a concept specific to agentic systems. Testing presents a key evidentiary gap for higher-tier use cases: the IIA Framework confirms that testing occurred but offers no technique for red teaming or hallucination testing. Third-party risk management provides a sound foundation that should be extended to agent permissions and change-notification testing.

CLOSING THE GAP

Building the translation layer

Taken together, the comparison yields a consistent conclusion: the OCC Framework defines the expected outcome, the IIA Framework supplies a methodology, and neither alone is sufficient. Institutions seeking to demonstrate examination readiness should construct a translation layer that converts each OCC expectation into an executable audit procedure with defined evidence.



In practice, that work includes

1. Map each in-scope SI use case to the relevant OCC components and assign a risk tier based on potential consequences, transfer of agency, and the degree of human oversight.
2. Identify where existing IIA-based procedures transfer directly (documentation and auditability, governance, internal audit independence) and where they require targeted extension (cybersecurity, data governance, third-party risk).
3. Develop bespoke procedures where neither framework offers ready-made techniques, including operational resilience, the quality of human oversight, memorization and inference testing, and hallucination-rate testing, drawing on the NIST AI Risk Management Framework and other frameworks and guidance where appropriate.
4. Add reconciliation testing that ties management and board reporting to the AI inventory, incident records, and audit findings.
5. Preserve the IIA's core principles of traceability, independence, and governance throughout, irrespective of the techniques selected.

READINESS CONSIDERATIONS

Questions every bank should be prepared to answer

- [ ] Does the organization maintain an inventory of AI use cases, tiered by potential consequence?
- [ ] How is human review of AI output designed and evidenced, particularly for customer-facing and other consequential decisions?
- [ ] How is AI functionality embedded in vendor products identified and overseen?
- [ ] How would an AI failure be contained and remediated, including for affected customers?
- [ ] Are the organization's testing and documentation sufficient to demonstrate that controls operate as represented?

How Helios can help

Helios' Risk & Audit practice supports banks and their internal audit functions in designing AI governance programs, building audit methodologies for SI use cases, and preparing for examination. An expanded version of the mapping summarized in this white paper is available upon request. To discuss how these frameworks apply to your institution, please contact Katie Granville or visit www.heliosgrc.com.

ABOUT THE AUTHOR

Katie Granville, Senior Partner

Katie Granville is a Senior Partner at Helios, where she leads the firm's risk and audit practice, including its AI governance, risk, and audit advisory work. She brings more than two decades of internal audit, risk, and controls experience at leading asset managers and financial institutions. Before joining Helios, Katie served as Chief Auditor of Allspring Global Investments (formerly Wells Fargo Asset Management) and previously led internal audit at MFS Investment Management. She has also held senior risk, controls, and compliance roles at Wells Fargo, Citizens Financial Group, and State Street, and serves on the AI Advisory Council of OnCorps AI. Katie advises banks, investment advisers, and asset managers on internal audit, risk management, the governance of emerging technologies, integrated data strategies, and AI enablement.

HELIOS

UNLOCKING NOW. DELIVERING NEXT.®

Helios provides governance, risk, and compliance advisory services to financial institutions and investment management firms, spanning compliance, legal, risk and audit, and cyber advisory.

heliosgrc.com

This white paper is provided for general informational purposes only and does not constitute legal, regulatory, or audit advice. Readers should consult qualified advisers regarding their particular circumstances.

© 2026 Helios. All rights reserved.