COSO releases GenAI roadmap

Internal control and GenAI

May 1, 2026

COSO releases new roadmap, Achieving Effective Internal Control Over GenAI.

Source and applicability

- COSO Roadmap: Achieving Effective Internal Control over Generative AI (GenAI)
- Applicable to all entities

Fast facts, impacts, actions

On February 23, 2026, the Committee of Sponsoring Organizations of the Treadway Commission (COSO) released a new roadmap, *Achieving Effective Internal Control Over Generative AI (GenAI)* (the roadmap), to guide entities in applying the 2013 COSO framework to GenAI.

The roadmap features include:

Tailored considerations	A capability-based approach
Provides GenAI-specific practices for each of the 17 principles in the COSO 2013 Internal Control – Integrated Framework.	Organizes GenAI use into eight distinct GenAI capabilities. For each capability type, illustrative focus areas and examples are presented across the five COSO components, with overarching control considerations by COSO principle.

Understanding the roadmap

Understanding how this roadmap fits within existing internal control frameworks and regulatory expectations is an important first step to applying its guidance appropriately. In particular, it is helpful to distinguish between COSO frameworks – developed through a formal issuance process to establish authoritative control expectations – and COSO-issued guidance intended to support the application of those frameworks in specific contexts, such as GenAI.

What the roadmap is and what it is not	
☑ COSO-issued guidance to assist in understanding how GenAI fits into the internal control considerations under the COSO 2013 Internal Control - Integrated Framework ☑ Real-world examples of GenAI uses and related control considerations ☑ Intended to be applied and tailored to an entity based on entity-specific facts and circumstances	☒ A formally issued control framework or update to the COSO 2013 framework ☒ A comprehensive control framework intended to establish required governance or control expectations ☒ A complete population of risks or controls

Interaction with the KPMG Guide

COSO's recently issued roadmap and the KPMG Guide, AI and automation in financial reporting (November 2024) address artificial intelligence (AI) risk from different but complementary perspectives.

The KPMG guide focuses on the AI solution level across the full AI lifecycle, which is the foundation from which multiple individual GenAI capabilities are deployed. It provides internal control over financial reporting (ICFR)-oriented guidance on how entities design, deploy and use AI solutions. Its scope extends beyond GenAI, recognizing that entities often combine multiple AI capabilities to achieve an intended objective. Further, it provides practical AI considerations for essential activities like understanding business processes, assessing information risks, and assessing Entity-level, General IT and Process control activities.

COSO's roadmap assesses risks and controls at the AI capability level, recognizing that GenAI risks may differ significantly based on the capabilities used. It includes examples at the capability level that may assist entities in implementing specific AI capabilities. The roadmap discusses principles around the emerging topic of 'AI reliance', providing a definition of reliance and, given the evolving nature of this emergent area, focuses on *what* is necessary for a reliable state, rather than prescribing *how* that reliability is achieved. The examples in the roadmap include illustrative metrics and artifacts – such as confidence scores, accuracy dashboards, or logs – which may assist entities in assessing the operation of AI. Such output will need to be validated through evolving testing approaches such as those listed in the roadmap (e.g. performance testing, multi-model validation, data analytics monitoring and third-party validation).

While the two documents have different objectives and distinct perspectives, together they provide a more holistic view of AI internal control considerations.

General reminders for using the roadmap

Reminder	What to consider	Why it matters
Understand all AI capabilities in use	AI solutions often incorporate multiple capabilities across a process, and those capabilities may be addressed by a single control activity. AI capabilities may also be leveraged within the entity's controls. Consider all capabilities in use.	Focusing on only one capability used within a process can result in incomplete risk identification and control response. Additionally, when a single control embeds multiple capabilities, clearly defining the ownership structure is essential to the effective assessment of its overall operating effectiveness.
Evaluate risks and controls end-to-end	When assessing risks and controls by capability, step back and review the process end-to-end to confirm that all relevant risks and controls have been identified.	Assessments focusing only on single or combined capabilities can inherently miss other critical aspects – e.g. the integrity of information during handoff that only become apparent at the full process level.
Evaluate the risks and controls at the relevant levels	Risks and controls may operate at different levels, including entity-level controls, process control activities and general IT controls (GITC). Consider whether risks are identified and assessed at the appropriate level so that the level of precision is commensurate with the risk.	Evaluating risks and controls at the appropriate level allows control design and precision to align with the nature of the risk and supports appropriate conclusions. Given the dynamic nature of AI, risks may arise at specific process integration points and may not be sufficiently mitigated by entity-level controls alone.
Determine ICFR implications of identified issues	When issues related to the design or output of AI are identified, in addition to implementing enhancements to the AI or surrounding processes, consider whether those issues also indicate control deficiencies.	When AI is involved in control activities, issues that result in changes to the AI's design or related controls may indicate ICFR implications that require evaluation and response.

Reminder	What to consider	Why it matters
Apply professional judgment to tailor and supplement as applicable	Information in the roadmap is intended to be informative and directional rather than exhaustive. It is meant to be applied and supplemented based on an entity's specific facts and circumstances.	Without professional judgment, guidance may not appropriately reflect entity-specific risks, controls and circumstances.

Contributing authors
Tommy Golden, Denae Hajovsky and Keith Hooper

Learn about us: **kpmg.com**

KPMG Financial Reporting View
kpmg.com/us/frv

This newsletter is part of our Defining Issues® collection of newsletters and articles with insights and news about financial reporting and regulatory developments.

Sign up here to receive news and insights delivered to your mailbox.