



TAXSERACT, LLC

U.S. & CROSS-BORDER TAX ADVISORY AND CONSULTING

taxseract.com · info@taxseract.com

August 7, 2026

Office of the Secretary
Public Company Accounting Oversight Board
1666 K Street, NW
Washington, DC 20006-2803

Re: Request for Public Comment on PCAOB Standard Setting (June 23, 2026)

Dear Members of the Board:

I appreciate the opportunity to comment on the Board's standard-setting and research agendas, and I commend the Board for opening its prioritization process to public input. I write as a certified public accountant with approximately twenty years of international tax and ASC 740 experience serving multinational issuers, most recently in Big Four practice, and as the founder of a technology company developing deterministic, traceable computation methods for tax and accounting workflows. That commercial interest is relevant to Recommendation V below and is disclosed there so the Board may weigh it.

This letter recommends that the Board's standard-setting and research agendas address seven areas in which observable, publicly disclosed developments in the artificial intelligence sector of the economy are placing novel demands on existing auditing standards. Each recommendation is grounded in disclosed transactions and published facts; none depends on any view of future market outcomes. The premise throughout is a modest one: auditing standards function as infrastructure, and infrastructure is most valuable when it is built before it is needed. Nothing in this letter should be read as implying any deficiency in the financial reporting of any named or identifiable registrant, or in the work of any registered firm; publicly reported transactions are referenced solely as evidence that the transaction classes discussed exist at material scale.

Summary of recommendations. First, that the Board undertake a research project on auditing interdependent financing and commercial arrangements — structures in which a vendor finances, invests in, or guarantees the obligations of entities that are simultaneously its customers — as they now appear at scale in the AI infrastructure supply chain. Second, that the Board's planned work on accounting estimates encompass the estimation-intensive revenue and purchase-price-allocation judgments characteristic of outcome-based AI contracts and AI-related acquisitions. Third, that the same estimates workstream address the valuation and disclosure of data and model assets generated by AI systems — including synthetic training data whose production is self-referential — where fair values rest on Level 3 inputs and internal reference prices. Fourth, that this workstream also address income tax valuation allowances whose realizability assessments depend on rapidly obsolescing technology assets. Fifth, that the Board's research agenda on technology include the evidentiary status of verifiable machine-generated records — deterministic logs and cryptographically verifiable computational lineage — as audit



evidence under AS 1105. Sixth, that the research agenda include engagement with the development, through broad multi-organization collaboration on the model of the COSO framework's origin, of suitable criteria for the assurance of AI systems whose outputs are material to financial reporting. Seventh, that the Board's estimates workstream and inspection program give focused attention to the audit of asset retirement obligation estimates for power-generation and data-center infrastructure whose continued operation depends materially on AI-related demand, particularly where retirement activities carry environmental or radiological risk.

I. Auditing Interdependent Financing and Commercial Arrangements in the AI Infrastructure Sector

Publicly disclosed transactions in the AI infrastructure sector now exhibit a recurring geometry: a supplier makes equity investments in, or extends financing or guarantees to, entities that are simultaneously among its largest customers; those customers finance purchases of the supplier's products with debt secured by the products themselves; special-purpose vehicles hold the underlying facilities; and the parties disclose multi-year purchase commitments and order books of unprecedented magnitude — in one widely reported instance, a stated one trillion dollars in purchase orders across two product generations. These arrangements are lawful, commercially rational, and publicly acknowledged. They are cited here for one reason: transaction classes of this geometry concentrate several existing auditing requirements onto judgments that current standards address only separately, and history teaches that structures in which the same economic interests occupy multiple positions in a revenue chain deserve early standard-setting attention. The Board itself was created in the aftermath of reporting failures involving round-trip and vendor-financed transaction structures; the auditing questions those structures raised were not novel individually, but their combination outran the guidance of the day.

The specific questions a research project should examine include the following. Under ASC 606, when a seller has provided financing to, or holds an investment in, its customer, the collectibility assessment and the identification of consideration payable to a customer both require judgments about the substance of intertwined flows — including whether, and when, amounts invested in a customer are in substance concessions that should reduce the transaction price. Under AS 2410, the auditor's evaluation of related parties and significant unusual transactions presumes that relationships of influence can be identified and their terms evaluated; where commercial dependence, equity participation, and financing run in a loop among formally unrelated parties, the standard's framework strains to describe relationships that are economically significant but fall outside its definitions. Consolidation and variable-interest-entity judgments concerning special-purpose facility vehicles determine whose balance sheet carries the assets and obligations of the buildout. Disclosure of remaining performance obligations and purchase commitments raises a sufficiency question where a material portion of a registrant's disclosed order book consists of commitments from counterparties in which the registrant holds interests or to which it has extended support. And going-concern evaluation



under AS 2415 presumes an entity whose prospects can be assessed substantially on its own; where a lattice of entities is mutually dependent — each entity's obligations serviceable only if the others perform — the auditor's unit of assessment deserves examination. I do not suggest that existing standards cannot reach these questions; I suggest that they could be reaching them piecemeal, and that a coordinated research project — potentially informing the Board's existing substantive analytical procedures, estimates, and related-party workstreams — would be timely while the transaction class and magnitude are still phasing in.

II. Auditing Accounting Estimates in Outcome-Based Revenue Arrangements and AI-Related Business Combinations

The Board has indicated that accounting estimates remain an area of standard-setting attention, and I recommend that this workstream explicitly encompass two estimate families now characteristic of the AI sector. The first is revenue recognition for outcome-based and consumption-based contracts. A growing class of AI service arrangements prices on results — per resolved case, per completed task, per measured improvement — or on consumption whose volume the customer controls. Under ASC 606 these features are variable consideration, and the registrant's estimate of the transaction price, application of the constraint, and allocation across performance obligations become the substance of revenue recognition. AS 2501 as modernized provides the auditor's framework; what it does not yet contemplate is estimation whose inputs are themselves machine-measured outcomes (was the case resolved? did the model's output meet the contractual threshold?) generated by the very systems whose performance is being priced. The auditor's evaluation of the relevance and reliability of such data as evidence, and of management's controls over the measurement systems, is an area where inspection-informed guidance would materially aid consistency.

The second family is purchase-price allocation in AI-related business combinations. Publicly disclosed transactions illustrate the estimation difficulty without any implication about the reporting entities involved: a disclosed \$587 million cash acquisition of an approximately sixteen-person AI startup, and a reported structure of approximately twenty billion dollars combining a technology license, an asset purchase, and the hiring of the seller's leadership and engineering team. Transactions of these classes require allocation among acquired technology whose economic life may be short and difficult to support, assembled workforce, contractual arrangements, contingent consideration measured at fair value with remeasurement through earnings, and — where structures combine license, asset, and employment elements — threshold judgments about what was acquired at all. Each allocation is an estimate; the observable characteristics of the sector (small target teams commanding large consideration; intangibles whose value derives from technology subject to rapid succession) concentrate unusual measurement uncertainty in these estimates. Guidance or staff practice alerts addressing the audit of intangible-asset valuations where the asset class exhibits documented rapid obsolescence would serve both auditors and investors.



III. Auditing the Measurement and Disclosure of Data and Model Assets Generated by AI Systems

A distinct measurement problem is emerging around synthetic data: data generated by AI systems — including the class of systems publicly marketed as “world models” — for the purpose of training, testing, and certifying other AI systems. Its defining economic feature, observable in current product offerings, is self-referential production: the systems that consume the data are of the same class as the systems that produce it, successive model generations are trained on data their predecessors generated, and within an enterprise or an interdependent group the producer and the consumer of the asset may be the same economic interest. Questions of recognition and measurement under the financial reporting framework are, of course, matters for the FASB, and this letter takes no position on them. The Board's interest arises because auditors already confront the asset class under existing standards, and each point concentrates unusual measurement uncertainty.

In business combinations, acquired data corpora and trained model weights are measured at fair value — routinely Level 3, with no observable market, with value derived from projections of internal use, and with economic lives subject to the documented obsolescence dynamics discussed in Sections II and IV. In capitalization judgments, the costs of generating data assets enter the financial statements through frameworks written for other things — internal-use software, film and content cost models — and the boundary between deductible operation and capitalizable asset creation is an estimate in itself. In impairment testing, recoverability of capitalized data and model assets depends on the succession of model generations, an input the registrant itself controls and documents. And in the evaluation of disclosure, a registrant's most economically significant productive asset may be one the framework does not recognize at all — such that the adequacy of financial statement and related disclosures about dependence on data assets, and about the *rights encumbrances* attached to them (training consents, licensing restrictions, and provenance limitations that directly bear on the rights-and-obligations assertion), becomes an audit judgment of first importance. Where the producing and consuming systems sit within one enterprise or among the interdependent parties described in Section I, management's valuations may rest on internal or circular reference prices; the auditor's evaluation of such evidence under AS 1105, and of the resulting estimates under AS 2501, would benefit from focused attention. I recommend that the estimates workstream explicitly include Level 3 valuations of data and model assets, and that the Board's staff coordinate observations with the FASB through existing liaison channels as that framework's own software-cost modernization and continuing intangibles research proceed.

One methodological observation may assist that workstream. Fair value measurement practice for intangibles leans heavily on market convention — comparable transactions, industry rules of thumb, customary discount and royalty benchmarks — and the asset classes discussed in this letter are precisely where such conventions are thinnest in data, literature, and practice. A substantial body of valuation methodology developed largely outside the financial reporting



framework addresses exactly this condition: the transfer pricing discipline, whose intellectual roots lie in management accounting's internal-pricing problem, has spent decades — under regulatory frameworks and in extensive litigation — developing structured approaches to valuing intangibles for which no market exists, including disciplined method-selection criteria, functional decomposition of value creation among contributors, evaluation of realistic alternatives, and contemporaneous documentation standards. I do not suggest that these approaches yield greater precision; the public record of transfer pricing disputes demonstrates otherwise. I suggest that they embody *process* rigor — explicit criteria for selecting among methods, and documentation disciplines that make valuation reasoning auditable — that could usefully inform both the evaluation of Level 3 estimates under AS 2501 and any staff guidance on auditing valuations of novel intangible classes. The convergence is already present in practice, where the same valuation specialists frequently serve both functions; the research opportunity is to examine what the audit of fair value estimates can systematically adopt from a discipline built for valuation in the absence of markets. Transaction classes of the kinds discussed in this letter are unlikely to remain confined to the current sector; structurally similar novel transactions can be expected to recur in adjacent frontier sectors, including the space economy, and methodological groundwork laid now will compound in usefulness.

IV. Auditing Income Tax Valuation Allowances Dependent on Rapidly Obsolescing Technology

A related but distinct estimate deserves separate mention because its difficulty is compounding: the valuation allowance assessment under ASC 740 for deferred tax assets — principally net operating loss carryforwards — of AI-sector registrants and registrants with significant AI-sector exposure. The realizability judgment requires scheduling future taxable income and the reversal of deferred tax liabilities. In this sector, the objective evidence bearing on that judgment has unusual features. Operating losses at technology developers can be large and sustained while the enterprise is, by design, in an investment phase. The economic life of the core productive intangibles is short and — unusually — *documented by the industry itself*: model developers publish depreciation and retirement schedules for their own systems, providing citable, contemporaneous evidence that successive generations of the principal asset class are withdrawn from service within periods measured in months to a few years. Deferred tax liabilities arising from accelerated recovery of hardware similarly reverse on the cadence of published product-generation cycles. None of these observations implies any particular conclusion for any registrant; collectively they mean that the positive and negative evidence framework of ASC 740, and the auditor's evaluation of it under AS 2501, must be applied to realizability horizons that are shorter, better documented, and more sensitive to technological succession than the framework's traditional applications. Inspection focus and staff guidance on auditing valuation allowance judgments in this fact pattern — including the auditor's use of the registrant's own published product-lifecycle evidence — would be timely.

V. Research Agenda: Verifiable Machine-Generated Records as Audit Evidence



Finally, the Board's request follows its earlier solicitation of views on the implications of artificial intelligence for financial reporting and the audit, and I offer a recommendation directed at the evidence layer rather than the analytics layer. Disclosure: my firm develops deterministic and traceable computation methods for tax and accounting applications, and I therefore have a commercial interest in the subject of this recommendation; I raise it because the underlying architectures are open, non-proprietary, and in production use across other trust domains.

Modern computational systems can be built to produce complete, tamper-evident records of what they did: deterministic logs of inputs, versions, parameters, and outputs, committed to append-only structures whose integrity is cryptographically verifiable. Architectures of this class are proven at scale outside financial reporting — certificate-transparency logging secures the web's identity infrastructure, and open provenance standards for digital content are in commercial deployment. As registrants increasingly generate the amounts and disclosures in their financial statements, and the data supporting them, through automated and AI-assisted systems, the auditor's foundational questions under AS 1105 — the relevance and reliability of information produced by the entity, and the controls over its production — can be answered categorically better where the producing systems emit verifiable records than where they do not. I recommend that the Board's research agenda include a project examining the evidentiary status of verifiable computational records: whether and how existing standards accommodate cryptographically verifiable system-generated evidence; whether such records, where available, bear on the auditor's assessment of the reliability of entity-produced information; and whether the standards should recognize, without mandating, the higher evidential quality of records whose integrity can be independently verified. The precedent is familiar: the standards already distinguish evidence by source and by the controls over its preparation. Extending that architecture to computational provenance would prepare the evidence framework for the systems that now produce the numbers — and would do so through the Board's ordinary research process, informed by the full range of stakeholders.

VI. Research Agenda: A Conceptual Foundation for AI System Assurance, Developed Through Broad Collaboration

The recommendation in Section V concerns the evidence layer; this final recommendation concerns the criteria layer above it, and it follows a precedent the Board knows intimately. The auditing of internal control over financial reporting does not rest on a framework the Board wrote: it operates against suitable, available criteria produced by a private collaboration of five professional organizations, formed in response to the reporting crises of its era through the Treadway Commission. The framework approach solved a problem structurally identical to the one now presented by AI systems whose outputs are material to financial reporting: how to give auditors an evaluable, common yardstick for complex entity-specific systems, without the standard-setter itself having to specify system design.



AI systems now generate amounts, estimates, and data that enter issuers' financial statements. What does not yet exist is a conceptual foundation for their assurance: a taxonomy of the classes of AI-system deficiency that are, in substance, material to financial statement users — accuracy and error behavior, drift, data lineage and rights, change control over models and training data, human oversight and override — together with severity classifications analogous to the deficiency/material-weakness hierarchy, and benchmark and scorecard instruments capable of serving as evidentiary matter. Critically, the criteria-based examination model the profession already operates demonstrates that such assurance can be delivered while fully protecting intellectual property and trade secrets: system-level examinations convey trustworthiness conclusions about systems whose internals remain confidential, and nothing about AI assurance requires disclosure of model weights, architectures, or training methods.

Components of such a foundation already exist in fragments: risk-management frameworks and management-system standards for AI published by national and international standards bodies, professional guidance emerging from the audit and information-systems assurance communities, and deep technical assessment literatures in the engineering and computing societies. What is missing is the integrating collaboration — spanning the accounting standard setters and their international counterparts, the audit and internal-audit professions, the information-systems assurance community, the engineering and computing societies, and the privacy and legal disciplines — that would do for AI system assurance what the Treadway-era collaboration did for internal control. I do not suggest the Board convene or lead such an effort; its independence and remit counsel otherwise. I recommend that the Board's research agenda include a project to (a) monitor and engage with criteria development across these communities, (b) identify what characteristics the Board would require before recognizing criteria as suitable for evidentiary purposes in audits of issuers whose reporting depends materially on AI systems, and (c) sequence any eventual standard-setting so that it can incorporate, rather than duplicate, a mature external framework. The scale of the collaboration contemplated is admittedly large; the confidence of the U.S. capital markets in financial statements increasingly produced by such systems is the corresponding stake, and the profession's own history shows the pattern succeeds.

VII. Auditing Asset Retirement Obligations for AI-Linked Power Generation and Data Center Infrastructure

A related estimate deserves separate attention because it sits on the retirement side of the same infrastructure buildout discussed above. Asset retirement obligations for data center facilities are already a recognized item under ASC 410-20 in filings across the sector, and nuclear decommissioning — recognized progressively as the facility is operated — has been treated as a paradigm example of a legal asset retirement obligation since the standard's origin. What is new is scale and provenance: aggregate capital committed to AI-related compute, data centers, and associated power generation is now estimated in the trillions of dollars over the balance of this decade, and a portion of the associated power sourcing is nuclear capacity, including at least one generating unit returned to commercial service under a long-term power purchase agreement



TAXSERACT, LLC

U.S. & CROSS-BORDER TAX ADVISORY AND CONSULTING

taxseract.com · info@taxseract.com

with an AI-infrastructure buyer several years earlier, then having been removed from service for economic reasons in a down turn. Facilities of this kind carry an asset retirement obligation whose funding adequacy rests on assumptions about the plant's remaining operating life and about the durability of the demand that justifies keeping it, or returning it, to service.

Two features of the current buildout strain those assumptions in ways the standard's traditional applications do not confront, and the Board's own regulatory counterpart has already identified both features as sources of historical funding shortfalls, even before AI-linked demand entered the picture. First, decommissioning cost estimates have in practice proven revisable as site conditions become better understood over the life of a facility, a documented pattern rather than a hypothetical one. Second, and more particular to the current buildout, a reactor brought into or kept in service principally to serve a bilateral power purchase agreement with a technology company operates as a merchant asset outside traditional cost-of-service rate regulation; unlike a rate-regulated utility reactor, such a facility has no assured recovery mechanism if the commercial relationship funding its operation ends, and the same regulatory history that documents the Connecticut Yankee cost overrun separately documents funding shortfalls at merchant reactors whose revenue depended on a single counterparty relationship. I do not suggest any parallel between the AI sector's demand drivers and any particular prior episode; investment decisions in nuclear generating capacity driven by regional demand projections have, at least once before and at comparable dollar scale, diverged materially from the assumptions underlying them, in an episode remembered as the largest municipal bond default in U.S. history at the time. That episode is offered solely to establish that decisions of this class and magnitude have previously proven wrong; it is not offered as, and should not be read as, a prediction about the current buildout, whose financing structures, counterparties, and demand drivers differ in kind.

I recommend that the Board's estimates workstream and inspection program give focused attention to the audit of asset retirement obligation estimates for power-generation and data-center facilities whose continued operation depends materially on a concentrated AI-related demand source, with three elements. First, that auditors evaluate the funding-adequacy and cost assumptions underlying such estimates on an interim basis, rather than only at year end, where the operating-life assumption rests materially on a single counterparty relationship or where retirement activities carry environmental or radiological dimensions — a cadence closer to the reporting frequency the Board's regulatory counterpart already requires of reactor licensees nearing the end of their projected operating life, applied here on a forward-looking rather than end-of-life-triggered basis. Second, that resulting guidance address disclosure sufficiency for funding adequacy specifically, extending to the retirement side of the buildout the same disclosure-adequacy principle Recommendation I raises on the acquisition side. Third, that the verifiable-record architecture recommended in Recommendation V above apply directly here: an interim-recertification expectation is only as reliable as the auditor's ability to confirm, after the fact, that it occurred on schedule and against the assumptions claimed. The Board's inspection program is again positioned to observe what no single engagement team can: obligations of this class are being incurred, and in at least one observable instance re-incurred after a period of



TAXSERACT, LLC

U.S. & CROSS-BORDER TAX ADVISORY AND CONSULTING

taxseract.com · info@taxseract.com

dormancy, across multiple registrants and facilities, at the moment their underlying demand and technology assumptions are least settled.

Closing

Each recommendation above concerns transaction classes and system characteristics that are observable today in public disclosures; none requires the Board to adopt any view about the future path of the sector. The common thread is timing: the confidence of the U.S. capital markets rests in part on audit and reporting standards that meet novel transaction classes early, and standards that do so have historically served as reference frameworks well beyond our borders. If the Board or staff would find it useful, I would welcome the opportunity to discuss any of these topics further.

Respectfully submitted,

Jason Wu, CPA, EA
Founder/Tax Principal,
Taxseract, LLC