

AI Governance at CalSTRS:

Summary of Findings from the Board-Staff Working Group on Technology Governance

May 28, 2026

Executive Summary

In October 2025, CalSTRS convened a Board-Staff Working Group on Technology Governance (the “Working Group”) to clarify the board’s responsibilities for overseeing technology at CalSTRS, with a particular focus on artificial intelligence (AI). The Working Group was formed in recognition that AI is no longer a future or experimental issue. AI is fast becoming a foundational layer of modern technology. Its use is increasingly embedded by default in standard systems and vendor offerings, and opting out may become more difficult over time. Already, AI influences internal operations, vendor relationships, enterprise risk, and the broader operating environment, and this influence will only expand as adoption becomes more widespread.

Through engagement with public pension peers, academic experts, and experienced advisors, the Working Group developed a clearer view of what effective board-level technology and AI governance requires. A consistent theme across these conversations was that AI should not be governed as a standalone initiative or technical program. Instead, it should be integrated into an organization’s existing strategy, risk management, and oversight practices. Boards are most effective when they focus on setting direction, establishing boundaries, and ensuring accountability; public pension boards like CalSTRS’ should not attempt to manage technology or its implementation.

The Working Group concluded that CalSTRS’ board already oversees many of the areas where AI introduces material implications, including long-term strategy, investment oversight, risk management, funding sustainability, organizational leadership, and member and public trust. The opportunity is not to expand the board’s scope, but to clarify how AI considerations should be embedded within existing governance responsibilities and decision processes.

Additionally, the peers and experts the Working Group spoke with reinforced CalSTRS’ global reputation for good governance. CalSTRS is recognized for having mature, robust governance processes relative to technology, risk management, and operations at both the staff and board levels. The opportunity presented in this report is for staff and the board to integrate AI governance into these existing practices.

Based on these learnings, the Working Group identified two initial governance-level actions for the board to undertake, supported by a small set of enabling recommendations related to integration, oversight, and external expertise.

- 1. The Working Group recommends that the board develop AI governance principles.**

CalSTRS currently has a staff-level operational policy governing AI use.

Governance-level principles would complement that policy by providing clarity at the board level on the use of artificially intelligent technologies in decision-making and organizational support. These principles are intended to play a clarifying and boundary-setting role, not an operational one. They will help the board and staff govern choices about AI over time, support consistent oversight, and provide a stable reference point as technologies, use cases, and risks evolve.

- 2. The Working Group recommends that the board develop a clear framework for overseeing AI-related risk.** Guided by the board’s AI governance principles, this

framework would support a risk- and materiality-based approach to oversight that informs when board involvement, enhanced oversight, or reliance on staff oversight is appropriate. This approach aligns oversight with risk and materiality, rather than novelty, and builds on CalSTRS' existing enterprise risk management practices.

In addition to these two initial actions, the Working Group offers enabling recommendations to strengthen the board's ongoing AI fluency.

1. **Ongoing Board Education:** Integrate AI-related topics more intentionally into the board's continuing education curriculum and onboarding processes.
2. **High-Level AI Use Overview & Readiness Assessment:** Request that staff provide the board a periodic, high-level report summarizing how AI is used across the organization and assessing the organization's readiness, framed to support the board's understanding of overall direction and risk awareness.
3. **External Expertise:** Take a deliberate approach to leveraging external expertise to provide the board with access to appropriate fiduciary, investment, governance, actuarial, and risk perspectives as AI-related issues arise.

Collectively, these recommendations ask the board to do three things: articulate board-level principles to guide the use of AI, confirm how AI-related risk should be overseen within existing structures, and direct how these responsibilities are integrated into current governance practices. These actions are intended to clarify board expectations without expanding the board's role into operational management.

From an implementation perspective, this report asks the CEO to integrate AI governance into existing board and committee workflows, provide the board with a high-level view of current AI use and organizational readiness, and support the board's development of principles and a risk framework. These actions are intended to reinforce disciplined governance while preserving appropriate flexibility and delegation as AI capabilities and adoption evolve.

Document Purpose

The purpose of this document is to summarize the Working Group's findings and to recommend next steps for the board to consider. The intent of this document is to describe what those next steps could involve and what governance questions may arise; the document is not intended to prescribe a particular outcome of these next steps.

Purpose and Scope of the Working Group

In October 2025, CalSTRS convened a Board-Staff Working Group on holistic technology governance to define the board's responsibilities in overseeing technology at CalSTRS, especially as it related to artificial intelligence (AI). The group was also tasked with developing perspectives on strategies for enhancing board fluency in technology and identifying relevant resources to support the board's evolving technology and AI governance practices.

Members of the Working Group included:

- Denise Bradford, Board Chair
- Harry Keiley, Investment Committee Chair
- Cassandra Lichnock, Chief Executive Officer
- Ashish Jain, Chief Technology Officer
- Dylan Pletcher, Chief Information Security Officer
- Scott Chan, Chief Investment Officer
- April Wilcox, Senior Investment Director

Trustees Michael Gunning (Board Governance Committee Chair) and Karen Yamamoto (Board Vice Chair) served as alternates. The board's governance consultant, Mosaic Governance Advisors, supported and advised the Working Group's efforts.

Between October 2025 and April 2026, the Working Group met with leading public pension organizations globally and engaged academic experts and external advisors. These discussions allowed the Working Group to explore approaches to technology and AI governance as well as to gather insights from peers that were farther ahead in their learning journey than CalSTRS. The scope of these engagements spanned pension administration and investment operations.

Recommendations for AI Governance Structures

The Working Group concluded that AI should be governed as an integrated part of CalSTRS' existing strategy, risk management, and oversight practices, rather than treated as a standalone technology initiative. Most board responsibilities related to AI can and should be embedded within existing policies, reviews, and decision processes. The board's role is not to manage AI, but to govern its use by setting direction, establishing boundaries, and ensuring appropriate delegation and accountability.

To enable the board to undertake its responsibility in a way that is integrated, disciplined, and consistent with existing governance practices, the Working Group recommends the following. These recommendations are organized to distinguish between actions that require board agreement and direction and actions that the board would direct the CEO to implement on the board's behalf.

1. Governance approach

The board should oversee AI at the board level and through its existing committee structures, as appropriate, by explicitly incorporating AI considerations into standing agendas, materials, and deliberations. The board should not create a separate technology or AI committee of the board.

2. Governance integration

The board should direct the CEO to develop and implement an approach for integrating the board's AI governance responsibilities into existing board practices, including governance processes, reporting, and decision-support materials, and to provide the board with periodic updates on progress, emerging issues, and any recommended adjustments.

3. Enhanced oversight of integration

The board should designate two trustees to work with the CEO – on a time-limited basis – to support the initial integration of AI governance responsibilities, help ensure clear and timely information flow to the board, and assist in preparing future board discussions and decision points related to AI.

4. External expertise to the board

The board should consider authorizing, where appropriate, the targeted use of external consulting resources to periodically assess and support the board's need for AI-related expertise as the technology, regulatory environment and risk landscape evolve.

These recommendations are intended to clarify governance responsibilities, reinforce appropriate delegation to the CEO, and provide a practical starting point for the board as it governs AI in an environment of ongoing technological change.

Core Board Responsibilities Related to AI

The board's responsibilities related to AI are not new or standalone responsibilities. Rather, they are an extension of the board's long-standing governance responsibilities, applied to an evolving technological context. Here we present core board responsibilities related to AI in a format that reflects how CalSTRS has traditionally understood and organized its board governance responsibilities. The board governs across four core domains – strategy, risk, performance, and people – threaded through the board's oversight of itself, the organization, and the plans, programs, and services it oversees.

I. Governance of the Organization

Strategy

- Clarify the role AI should play in advancing CalSTRS' long-term strategy and mission, values, and fiduciary obligations within applicable legal and regulatory requirements.
- Consider AI-related opportunities and risks as part of long-term strategic planning and scenario discussions, consistent with the board's oversight role.

Risk

- Establish high-level governance principles and a risk appetite for the use of AI that are consistent with CalSTRS' mission, values, and fiduciary obligations.
- Ensure AI-related risks – including legal, regulatory, privacy, data-protection, and accountability risks – as well as unintended consequences are identified, monitored, and escalated through existing risk governance and enterprise risk management (ERM) frameworks.

Performance

- Monitor whether significant AI initiatives are appropriately scoped, resourced, and aligned with board-approved strategic and budget priorities.
- Ensure that board-level reporting provides sufficient visibility into material AI uses and their implications, commensurate with risk and impact.

People

- Delegate accountability for enterprise-wide AI strategy, governance, and risk management to the CEO, consistent with existing authorities and accountability mechanisms.
- Ensure that leadership capacity, talent strategy, and change-management approaches appropriately consider the organizational impacts of AI adoption over time.

Existing governance mechanisms include: strategic planning, annual budget approval, ERM (including legal, regulatory, and compliance risk), executive reporting, CEO evaluation, and succession planning.

II. Governance of Plans, Programs, and Services

Strategy

- Periodically assess whether AI-driven changes in investment markets, service delivery, or the broader operating environment warrant consideration in board-level strategies, policies, or beliefs.

Risk

- Ensure that staff maintains appropriate processes for monitoring and managing AI use by external entities, including investment managers, consultants, vendors, and other business partners, with due consideration to contractual, regulatory, and compliance implications.
- Oversee how AI-related risks (including compliance and reputational risks) affecting the investment branch, member and employer services, or system sustainability are reflected within existing oversight and risk frameworks.
- Ensure the System's framework for administering plans, programs, and services in compliance with applicable law considers emerging AI-related laws and/or regulations.
- Ensure the System's framework for keeping member and organizational data secure includes provisions for the use of such data in artificially intelligent systems.

Performance

- Oversee how AI use affects outcomes that matter to the board, such as investment performance, quality and reliability of services, and long-term system sustainability.

People

- Recognize that AI use in plans, programs, and services may affect member, employer, and beneficiary experiences, and ensure that staff considers those impacts in design, deployment, and communication, consistent with the board's expectations and values.
- Rely on staff to address workforce, training, and operational matters related to AI-enabled services, consistent with established authorities and accountability structures.

Existing governance mechanisms include: investment beliefs and policies; risk reporting; vendor oversight, procurement and contracting practices; actuarial experience studies; funding risk management; customer service; and stakeholder reporting.

III. Governance of the Board

Strategy

- Ensure the board maintains sufficient fluency to engage meaningfully in governance-level discussions about and oversight of AI over time. Fluency is defined as the ability to ask informed questions and exercise judgment in setting direction and ensuring accountability; it does not require technical expertise.

Risk

- Establish clear expectations regarding trustees' use of AI tools in support of board work, consistent with fiduciary duties of care, loyalty, and prudence and with board policies regarding confidentiality, records, and public accountability.

Performance

- Periodically assess whether board structures, agendas, and information flows are effectively surfacing AI-related issues that warrant governance-level attention.

People

- Integrate AI into the board's continuing education and onboarding curriculum in a manner appropriate to the board's oversight role.
- Periodically review and refresh AI-related education to reflect evolving technologies, risks, and governance expectations.
- Periodically assess the need to engage external experts and specialists to support the board in continuing to build AI fluency.

Existing governance mechanisms include: Board education programs; board and trustee self-evaluation; governance policies; and committee charters.

Next Steps

The Technology Governance Working Group originally focused on emerging technologies more broadly but quickly realized that a specific focus on AI was needed given its transformative potential, widespread use, and disruptive nature.

Through its engagements with peers, academics and experts, the Working Group identified several practical next steps for the board to pursue related to its role in governing AI usage within the organization. These next steps are meant to clarify the board's guidance for the staff's use of AI, provide governance level direction that dovetails with existing operational policy, and to build greater fluency across the organization regarding AI's impact.

Initial Board Activities

The Working Group recommends that the board undertake two initial actions over the near term that will create greater clarity around the board's perspective on AI and its governance role in an AI-enabled environment. These actions will serve as a clear foundation for ongoing AI oversight. The Working Group also identified additional next steps that should be directed to the CEO to undertake with periodic updates to the board.

1. Board-Level AI Governance Principles

The Working Group recommends that the board develop AI governance principles.

These principles will serve a clarifying and boundary-setting role, not an operational one. Their purpose is to help the board and staff govern choices about AI, rather than manage the technology itself. In practice, governance-level AI principles are intended to:

- Set direction, not rules, by articulating the values and expectations that should guide the use of AI without prescribing tools, models, or controls.
- Ground decision-making over time, providing a stable reference point as technologies, vendors, and use cases evolve.
- Create shared expectations across leadership, supporting alignment between the board and staff on how AI should be used in service of the CalSTRS mission and fiduciary responsibilities.
- Enable oversight without micromanagement by allowing the board to assess consistency with principles rather than engaging in technical detail.

These governance-level AI principles provide longer-term direction that informs judgment and frames oversight conversations while leaving room for staff to adapt and be nimble as technology and risks change. These principles may evolve over time as AI capabilities and risks change and other intelligent technologies for decision-making emerge.

CalSTRS currently has an internal, staff-level operational policy that governs execution, controls and staff use of AI. Developing governance-level AI principles would complement that policy by providing completeness at the board level.

To inform the principles, the Working Group recommends that the board also discuss core beliefs about the definition of AI, AI's transformative potential, the risks AI poses, and other similar topics.

2. AI Risk Framework

The Working Group recommends that the board develop a clear framework for overseeing AI related risk.

AI introduces risk to the organization, both in ways addressed through current risk management practices and in new ways. The board needs a framework for identifying and assessing material risks related to AI based on risk characteristics, potential impact, and materiality, consistent with CalSTRS' existing enterprise risk management (ERM) approach. It also needs visibility into how AI initiatives are overseen across the organizational governance structure.

Using the AI governance principles it develops as guidance, the Working Group recommends that the board develop a clear risk framework that guides how the board further defines and acts on the responsibilities outlined in this document. The AI risk framework should be designed to integrate AI-related risks into existing ERM practices and could support the board in the following ways:

- Support consistent board-level judgment about where AI use may raise risks or implications that warrant board involvement, enhanced oversight, or the establishment of guardrails, versus where reliance on existing delegation and controls is appropriate.
- Enable differentiation among AI uses based on relative risk, materiality, and sensitivity.
- Inform the establishment and periodic review of a risk appetite that accounts for various types of risk introduced by AI, such as fairness and human-oversight risks, system reliability and data-quality risks, compliance and auditability risks, ecosystem or reputational risks, and workforce-related risks.
- Define an appropriate cadence for AI risk-related reporting and communication between the board and management.

This risk framework should be routinely reviewed and refined as needed to ensure it continues to align with CalSTRS' ERM framework and appropriately covers emerging and evolving risks.

Additional Next Steps

The two initial board actions will establish direction and boundaries. The Working Group also recommends additional next steps that would be advanced by staff, at the board's direction, over the course of the twelve to eighteen months.

These next steps are intended to inform and support the board's governance role, and they will provide information and context to support implementation of the board's initial actions. These items are designed to provide context, visibility, and decision-useful information rather than to expand board involvement in operational or technical matters.

1. Provide High-Level AI Use Overview and Staff Perspective

To support effective governance and informed discussion, the Working Group requests that staff periodically provide a board-level overview of current AI usage at CalSTRS.

The purpose of this overview is to provide trustees with clarity about where and how AI is used in ways that may be relevant to strategy, risk, performance, or people. The overview should be framed at a level appropriate for board oversight and should not introduce unnecessary technical detail. Rather than cataloging all tools or technologies in use, staff should focus on categories of AI use, the roles those uses play, and the areas of the organization where AI may have material implications. This overview should be provided periodically as AI usage will evolve over time.

This overview is not intended to imply board approval or endorsement of any particular tool, vendor, or application. Responsibility for selecting, implementing, and managing AI tools remains appropriately delegated to staff. A detailed, item-by-item inventory is unlikely to be useful for governance purposes and could create confusion or reputational risk if interpreted outside its intended context (i.e., a tool's inclusion in this overview should be understood as information provided to support governance judgment only).

In addition to describing current usage, the Working Group believes it is important for trustees to understand staff's perspective on how AI supports CalSTRS' mission and strategic objectives. Staff are therefore encouraged to describe how AI usage advances organizational goals, where AI is viewed as most relevant or impactful, and how AI use may evolve over time.

At a board-appropriate level, this overview could help trustees understand:

- The primary categories of AI use across the organization and their relative risks and impacts.
- How AI supports CalSTRS' long-term effectiveness, sustainability, and service delivery.
- How staff monitors emerging AI developments and reassesses potential opportunities and risks.

- Whether there is an overall approach to AI adoption and integration, recognizing that specific uses may continue to develop over time.

The intent of this request is not for the board to set AI strategy or direct individual use cases. Rather, it is to ensure that trustees have a clear and shared understanding of how staff use AI so that trustees can apply the board's long-standing governance responsibilities thoughtfully. This understanding provides essential context for the board's work to clarify principles, set risk tolerance, and distinguish where board involvement is required, where ongoing oversight is appropriate, and where reliance on existing delegation and controls is sufficient.

The Working Group anticipates that this overview could be provided in a concise briefing or dashboard format that highlights categories of use, staff's view of relative risk, and areas of impact, rather than detailed system or vendor inventories.

2. Assess Organizational Readiness for AI Use

The Working Group believes it would be helpful for staff to periodically provide a board-level overview of CalSTRS' organizational readiness as it relates to current and potential AI use.

In addition to understanding how AI is used, effective governance also requires some visibility into whether CalSTRS is organizationally prepared to support responsible AI use. AI capabilities are fundamentally dependent on the quality, availability, and governance of underlying data. As a result, data readiness is often a more limiting factor than the AI tools themselves. Additionally, AI use will impact the organization's workflows, processes, and resource needs.

This request for a periodic readiness overview is not intended to be a technical assessment or a review of data systems. Rather, it is meant to help trustees understand whether foundational conditions are in place to support AI use consistent with CalSTRS' strategy, risk tolerance, and fiduciary obligations.

This overview could address the following at a high level:

- Whether key data used for decision-making and service delivery is sufficiently reliable, governed, and accessible to support responsible AI use.
- How data classification, quality, privacy, and security considerations are addressed within existing governance and control frameworks.
- Known limitations or gaps in data readiness that may constrain AI use or elevate risk.
- Whether current or anticipated resource needs related to data governance, talent, or infrastructure are aligned with staff's expectations for AI use over time.

The intent of this discussion is not for the board to direct data strategy or technical decisions. It should provide sufficient context for understanding the pace and scope at which AI can be adopted responsibly to ensure that the board's AI-related principles, oversight expectations, and resource discussions are grounded in organizational realities.

This readiness perspective, alongside the AI usage overview, serves an expectation setting purpose. It will help the board distinguish between what is possible in theory and what is prudent in practice. These two assessments will support more informed judgment about where AI use may warrant board involvement, oversight, or reliance on existing delegation.

3. Enhance Board and Trustee Education

The Working Group recommends that AI be explicitly incorporated into the board's continuing education curriculum as a standing topic.

The board currently has a well-developed continuing education framework that encourages a mix of learning modalities, including external educational programs sponsored by third parties, in-boardroom education, and self-paced education. In addition, staff are available to meet with trustees to support their understanding of how AI is used within the organization and the internal governance processes that apply. The board has also received AI-related education sessions this year that were recorded and may be made available to trustees for ongoing reference.

As AI becomes more embedded in the organization's operations, systems, and processes, a more intentional and structured approach to board education is warranted. Effective AI governance does not require trustees to develop technical expertise, but it does require a shared baseline of understanding to support informed oversight, sound judgment, and appropriate delegation. Over time, trustees are likely to encounter AI-related considerations not as standalone topics but as part of routine board discussions and decisions related to investments, services, risk, and organizational performance.

Consistent with the board's governance role, AI education should prepare trustees to recognize when AI is relevant to a decision and understand its implications for fiduciary oversight. The objective is to support proportional and thoughtful governance, aligned with the board's long-standing responsibilities. Education resources do not need to equip trustees to evaluate technical design or implementation details.

The Working Group recommends the incorporation of AI as a standing topic within the board's continuing education curriculum. This approach would help ensure that trustees develop and maintain sufficient fluency to engage meaningfully on AI-related risks, opportunities, and tradeoffs as part of their fiduciary responsibilities.

The Board Governance Committee is responsible for oversight of the board education curriculum. Specifically, the Working Group recommends that the Board Governance Committee:

- Integrate AI- and technology-related topics into the continuing education curriculum, with content tailored to the board's oversight role rather than staff's operational responsibilities.

- Define a number of learning hours related to technology and AI to be completed within a trustee's first two years of service, recognizing the pace of change and the relevance of these topics to long-term governance.
- Incorporate the presentations delivered to the full board by Dr. Ashby Monk and Dr. Param Vir Singh into the new trustee onboarding curriculum as foundational learning resources for the foreseeable future.
- Ensure that trustees receive communications from staff regarding relevant, high-value technology-related training and educational opportunities.

Collectively, these steps will help ensure that board education on AI is proactive, consistent, and aligned with the AI governance principles the board develops while preserving the appropriate distinction between oversight and staff.

In addition, the Working Group recommends that AI-related education opportunities available to trustees be reviewed and refreshed annually. Given the rapid pace of technological change and the evolving risk and governance landscape, periodic reviews will help ensure that the board's continuing education remains relevant, appropriate for fiduciary oversight, and aligned with emerging practices and issues. This review should be overseen by the Board Governance Committee and informed by shifts in AI use within the organization and developments in the broader investment and regulatory environment. The committee's assessment of evolving trustee understanding may also uncover a need to adjust board agendas, information flows, or governance practices related to AI.

4. Board Use of Consulting Resources

The Working Group recommends that the board engage external experts to support the board's ongoing governance of AI.

The board's fiduciary responsibilities require it to have access to qualified expertise in areas where specialized knowledge is necessary to support informed oversight and sound judgment. As AI becomes more relevant to enterprise operations, risk management, and investment decision-making and operations, no single advisor or function can reasonably address all implications for the board.

Accordingly, the Working Group recommends that the board take a deliberate and coordinated approach to periodically assessing the need for and utilizing external expertise to support its governance of AI. This approach is consistent with the board's long-standing practice of engaging subject-matter experts to inform complex oversight responsibilities while preserving the appropriate distinction between board oversight and staff execution.

As an initial step, the Working Group recommends that the board engage an external expert to advise on the development of AI governance principles. External perspective can help ensure that the principles are appropriately scoped for board-level governance and informed by emerging practices.

Looking forward, the board should remain open to engaging external expertise as needed to support effective oversight of AI as use cases, risks, and regulatory expectations evolve. To that end, existing and future consulting resources should be applied intentionally and within clear governance boundaries. Examples of expertise that may be relevant, depending upon context, are as follows:

- AI governance expertise could support the board in establishing and periodically reviewing AI governance principles.
- Enterprise risk and other specialized consultants, as appropriate, could assist the board and organization in developing and/or validating an integrated approach to overseeing AI-related risks and aligning that oversight with the CalSTRS enterprise risk management framework.
- Fiduciary Counsel could assist the board in determining how the duties of loyalty, prudence, and care apply in the context of AI-enabled decision making. This could include advising on the fiduciary implications of AI use in operations, the board's appropriate reliance on AI-informed inputs, and other relevant legal and regulatory developments, including emerging international standards that may impact fiduciary expectations.
- Board Investment Consultants could support the Investment Committee's understanding of AI-related considerations within investment portfolio governance, including how AI may affect decision-making, risk management, and competitive positioning.
- Staff actuarial leadership and the board's external actuarial consultant could assist the board in understanding how AI considerations intersect with funding assumptions, modeling, and long-term funding risk management.
- The board's governance consultant could support the board, Board Governance Committee, and CEO in integrating AI considerations into existing governance structures – including the board and committee workplans, charters, policies, and practices – to reinforce a proactive governance posture.
- Other experts could be engaged as needed to address evolving risks and opportunities.

Overall, this approach enables the board to draw on specialized expertise where it adds the most value.

Report Conclusion

AI is evolving rapidly and is increasingly relevant to how organizations such as CalSTRS think about future strategy, operations, and risk. While CalSTRS has established operational policies and controls governing the use of AI, the continued development of the technology raises important governance considerations for the board. The Working Group was convened to help clarify the board's oversight role in this evolving environment consistent with CalSTRS' fiduciary obligations.

The Working Group's recommendations are intentionally focused and grounded in the board's existing governance framework. They are designed to strengthen how the board governs AI by emphasizing principles, integration into existing practices, and appropriate delegation without creating new committees or drifting outside its board governance lane.

The initial actions and enabling recommendations outlined in this report provide a practical starting point for the board to govern AI in a disciplined and role-aligned manner. Together, these steps support informed oversight, preserve flexibility as technologies evolve, and reinforce CalSTRS' long-term governance approach.

Appendix A. Key Learnings

Between October 2025 and April 2026, the group met with two leading public pension organizations globally and engaged two academic experts and two external advisors. These discussions allowed the group to explore approaches to technology and AI governance and gather best practices. The scope of these engagements spanned both pension administration and investment operations.

Working Group members Cassandra Lichnock and Amy McDuffee held additional conversations with several peers and experts. Peers that were less advanced in their AI journey or were unable to provide additional insight into their AI governance strategies were not advanced to a session with the full Working Group. Some peers declined due to scheduling challenges and/or limited capacity to engage. Only experts with experience related to board-level AI governance were invited to interact with the full Working Group.

Peer 1

This peer was interviewed by the Working Group in October 2025.

- The fund's initial emphasis on technology (including but not limited to AI) was driven by an urgent need to insource key capabilities for member services and a fundamental strategic shift toward internal management of the investment portfolio.
 - Additionally, national regulations have reinforced the need for specific board oversight of technology.
 - Subsequently, the board's overall strategy has led to continued prioritization of technology.
- The strategic decision to build internal technological capabilities led to the formation of a technology subcommittee. The subcommittee leveraged independent advisors and external specialists to support board oversight.
 - That subcommittee was disbanded as the organization matured. Today, technology is under full board oversight, and the fund uses a technology reference group of executive leaders to monitor technology strategy.
- To support this evolution, the fund developed a maturity scorecard with the support of an external consultant to assess progress. Once the fund reached its maturity score goal, the scorecard was simplified for ongoing oversight.
- Because the fund determined that technology was an important strategic lever, the board prioritized building technology fluency, which they did through academic programs and conversations with technology firms.
 - Management participates in these educational efforts to ensure a shared vocabulary and alignment.
- The fund created an innovation lab (with a separate, minimal budget) to pursue targeted exploration of methods for evolving the organization. Pilot projects must be tied to member benefits.

Peer 2

This peer was interviewed by the Working Group in December 2025.

- Technology governance is focused on creating alignment between organizational priorities and technology objectives
 - The governance is used to ensure technology investments generate business value, improve risk and resource management, and increase the organization's ability to achieve its goals.
- The organization leverages the COBIT framework, which is an industry-recognized and proven governance structure
- The organization has a distinct focus on AI, which is governed through a set of six high-level principles that serve as guardrails that create both freedom and accountability.
 - These principles were developed by management and approved by the board.
- Technology oversight is integrated with other board responsibilities. They do not use a separate committee.
 - AI risks are treated like any other risks – they are integrated into the Enterprise Risk Management system.
- The board oversees the business plan, which includes the organization's core ambitions. Technology strategy and oversight are a result of and directly tied to those ambitions.
 - Projects with budgets over a certain threshold are given board oversight committees.
- Innovation is embedded in the organization, with top-down and bottom-up mechanisms to continuously evolve the organization. The technology director ensures they are responsive to problems and where the organization needs to build new capabilities – new technologies are not the starting point.

Academic Expert 1 – Dr. Ashby Monk, Stanford University

Dr. Monk presented to the Working Group in November 2025 and later presented to the full board on January 28, 2026.

- Organizations need board-level clarity on AI governance, data foundations, and organizational readiness. The board's role is to articulate why the organization is investing in AI, what risks are acceptable, how success will be judged, and what resources will be committed.
 - Pension funds have three levers for change: governance, culture, and technology. Governance and culture are slow and difficult to change; technology offers a more immediate and feasible mechanism to drive organizational evolution.
 - Boards do not need to become technical experts, but they must understand enough to set direction, resource the organization appropriately, and establish guardrails for AI usage.

- The next generation of technology is defined by advancements in inferential depth (e.g., the ability to generate insights, understand problems differently, and identify patterns) not by advancements in speed.
 - For AI to deliver value, organizations need clean, reliable, and well-governed data.
- Dr. Monk used a GPS analogy to describe future AI-enabled systems: organizations need a precise location (current state), a clear destination (true objectives), and smart routing (strategies and scenarios). Most pension funds still operate with “paper maps” (e.g., spreadsheets and static dashboards) rather than dynamic, integrated positioning systems.
- AI adoption is primarily an organizational and cultural challenge, not a technical one. Early pilots should position AI as a parallel tool that checks work and analyzes past decisions rather than replacing current processes outright.
 - Meaningful progress may require a dedicated sandbox or lab with its own budget, safe data, and contracting flexibility. Incremental “baby steps” risk failure due to fear, internal resistance, and lack of organizational commitment.

Academic Expert 2 – Dr. Param Vir Singh, Carnegie Mellon University

Dr. Singh presented to the Working Group in January 2026 and later presented to the full board on March 5, 2026.

- Pilot projects most often fail because of organizational issues – not technical failures.
 - For these new solutions to fully integrate into an organization, they need to solve a real problem and have a clear value statement up front. Organization strategy should drive technology strategy.
 - Technology adoption is a change management challenge; it may require reorganization of existing processes and workflows.
- The board sets the boundaries within which AI can be implemented and is thus accountable for all AI-supported decisions. The board establishes accountability and defensibility of AI usage.
- The goal of building board fluency in technology matters is not to develop technical experts but to ensure trustees can ask the right questions, establish proper guardrails and decision-making processes, and create a monitoring system that proactively flags outliers.
 - Fluency also means knowing how to assess emerging risks and opportunities.
- The board should adjust its oversight to the risk and material impact of technology – not its novelty.
- Traditional ROI measures will not work for AI. Qualitative measures are needed to express AI’s value in the short-term, and the board should establish metrics that indicate trends toward value that can be used to monitor progress and make go-no-go decisions.
- AI risks are not like cybersecurity risks. The four major types of risk the board should look at are: (1) people and fairness; (2) systems and performance; (3) institutions and compliance; and (4) ecosystem and reputation.

External Advisor 1

An external advisory firm, referred to CalSTRS through Dr. Singh, presented to the Working Group in January 2026.

- AI is not “just another tool”; it represents a fundamental, economy-wide shift comparable to electricity. Boards must actively push management toward adaptation rather than assume it will happen organically. The pace of change means boards cannot wait for clarity; governance must evolve in parallel with uncertainty.
- The board’s role encompasses three core responsibilities:
 - Architects of aspiration: Setting ambition, stretching the organization’s vision, and ensuring AI is used where it truly matters.
 - Avid challengers: Asking uncomfortable questions early, pushing management on both opportunity and risk, and avoiding complacency.
 - Custodians of performance: Establishing guardrails, accountability, and oversight so experimentation does not create unacceptable downside.
- The half-life of technological expertise is extremely short; even highly specialized experts can become obsolete quickly. Rather than relying on individual trustees to bring expertise, boards should focus on raising their collective AI fluency and leverage outside experts.
- AI should be looked at from both opportunity and risk standpoints. Boards should use scenario planning and dynamic planning processes to imagine possible futures and their risks.
- Most organizations still govern AI on a case-by-case basis, but leading organizations are moving toward principles-driven governance aligned with strategy and regulatory context. These principles are valuable in creating a common language.
 - Establishing what is off-limits is often as important as defining what is allowed.

External Advisor 2

A regulatory expert, referred to CalSTRS through a peer, spoke with several Executive Staff members and then presented their findings to the Working Group in April 2026.

- The board’s role in AI governance is to set the boundaries within which management innovates, set risk appetite, ask strategic questions, and hold management accountable for answers. The board’s most important contribution is defining where innovation meets organizational values.
- AI is already embedded in the information systems and vendor relationships that CalSTRS depends on. The question is not whether to engage with AI but whether CalSTRS governs it deliberately or by default.
- AI governance is not a new discipline requiring new infrastructure. It should be part of existing enterprise risk management frameworks. It does, however, require a sharper focus on risk appetite, governance of AI in operations, and vendor due diligence.

- At its core, AI is software that learns from data and makes predictions or decisions without being explicitly programmed for every scenario. AI makes decisions at scale that previously required human judgement, and it often works in ways its creators cannot fully explain. This is why governance matters: you need guardrails when you cannot see exactly how decisions are made.
- Privacy, data protection, employment standards, and transparency laws all apply to AI systems right now. California is an active AI regulator, and there are several pieces of state legislation that apply to CalSTRS' use of AI. CalSTRS legal services actively monitors this landscape.
- The EU is furthest ahead in terms of regulating AI, with a comprehensive AI Act that categorizes AI harms by risk level and assigns responsibility proportionately. Compliance with the EU's act will shape product design globally. The EU's high-risk categories are also instructive as best practice.
- Adoption risk is real, but the risk of inaction is equally real. Implementing AI systems without adequate governance can embed biases, create legal liability, erode data integrity, and produce unintended or unexplainable outcomes. However, organizations that over-restrict AI adoption out of excessive caution will fall behind.
- Because CalSTRS is mainly a procurer of AI rather than a developer, its single largest governance exposure is vendor management.

Appendix B.

Board-Level AI Governance Matrix (Reference for Trustees)

How to use this tool: This matrix is intended to help trustees quickly orient themselves to their board-level responsibilities related to AI during board and committee discussions. AI does not introduce new fiduciary duties. It reinforces existing responsibilities across strategy, risk, performance, and people, applied across three areas of board oversight.

Board-Level AI Governance Matrix

Oversight Area	Strategy	Risk	Performance	People
Governance of the Organization	Clarify the role AI should play in advancing CalSTRS' long-term strategy and mission, values, and fiduciary obligations within applicable legal and regulatory requirements. Consider AI related opportunities and risks as part of long-term strategic planning and scenario discussions.	Establish high-level governance principles and a risk appetite for the use of AI. Ensure AI related risks as well as unintended consequences are identified, monitored, and escalated through existing risk governance and ERM frameworks.	Monitor whether significant AI initiatives are appropriately scoped, resourced, and aligned with board-approved strategic and budget priorities. Ensure board-level reporting provides visibility into material AI uses proportionate to risk and impact.	Delegate accountability for enterprise-wide AI strategy, governance, and risk management to the CEO. Ensure leadership capacity, talent strategy, and change management consider the organizational impacts of AI adoption.
Governance of Plans, Programs, and Services	Periodically assess whether AI driven changes in investment markets, service delivery, or the broader operating environment warrant consideration in board-level strategies, policies, or beliefs.	Ensure appropriate processes are in place to monitor and manage AI use by external entities with due consideration to contractual, regulatory, and compliance implications. Oversee how AI-related risks are reflected within existing oversight and risk frameworks (including the framework for administering plans, programs, and services in compliance with applicable law as well as the framework for keeping member and organizational data secure).	Oversee how AI use affects outcomes that matter to the board, such as investment performance, quality and reliability of services, and long-term system sustainability.	Recognize that AI use in plans, programs, and services may affect member, employer, and beneficiary experiences, and ensure that staff considers those impacts in design, deployment, and communication. Rely on staff for workforce, training, and operational decisions, consistent with delegation.
Governance of the Board	Maintain sufficient fluency to engage meaningfully in governance-level discussions about AI. <i>(Fluency means the ability to ask informed questions and exercise judgment.)</i>	Establish clear expectations regarding trustees' use of AI tools in support of board work, consistent with fiduciary duties of care, loyalty, and prudence and with board policies regarding confidentiality, records, and public accountability.	Periodically assess whether board structures, agendas, and information flows are effectively surfacing AI related issues that warrant governance level attention.	Integrate AI into the board's continuing education and onboarding curriculum. Periodically review and refresh AI-related education. Periodically assess the need to engage external experts.

Reminder for Trustees

When AI arises in a board or committee discussion, ask:

- **Which domain is this?** Strategy, risk, performance, or people?
- **Which oversight area does it affect?** The organization, plans/programs/services, or the board itself?
- **What is the board/committee's role here?**
 - Approval of approach/guardrails (highest-risk uses)
 - Oversight (material risks or impacts)
 - Reliance on delegation (lower-risk uses within established controls)