CITY OF SAN JOSE
CAPITAL OF SILICON VALLEY

Memorandum

TO: Board of Administration for the Federated City Employees' Retirement System

FROM: John Flynn

SUBJECT: AI Governance at San José and ORS:

DATE: September 17, 2026

RECOMMENDATION

Staff recommends the Board adopt the City of San José Artificial Intelligence Policy 1.7.12 as ORS's governance.

BACKGROUND

The Office of Retirement Services (ORS) administers the Federated City Employees' Retirement System and the Police and Fire Department Retirement Plan through two legally independent Boards of Administration. This memorandum accompanies a staff presentation requested by the Board and the Governance Committee chairs on Artificial Intelligence (AI) governance. Its purpose is threefold: to describe the City of San José's current artificial intelligence governance structure as it applies to ORS, to identify where ORS's current practice falls short of that structure, and to describe the internal governance process staff is establishing to close those gaps and to govern ORS's use of artificial intelligence going forward.

Artificial intelligence use across ORS today consists of staff productivity work, drafting, research, summarization, and document review, conducted with human review throughout. No ORS division has identified an instance in which artificial intelligence makes a final benefit, financial, investment, or management decision independently. The Investment Office is, however, beginning its first system in which artificial intelligence functions are embedded directly into a workflow, within its new EnvestHub platform, discussed further in the Analysis below. Given this transition, and at the request of the Board and the Governance Committee chairs, staff believes this is an ideal point to review the governing framework and ORS's functioning within it.

ANALYSIS

The City of San José's AI Governance Framework

The City governs AI through two related but distinct sets of requirements: one addressing the procurement and deployment of AI-enabled systems, and the other addressing individual staff use of GenAI tools.

1) **Procurement and Deployment of AI-Enabled Systems.** City of San José Artificial Intelligence Policy 1.7.12, together with the City's AI Handbook, governs the purchase, configuration, development, use, and maintenance of AI systems. The Policy states eight guiding principles: effectiveness, transparency, equity, accountability, human-centered design, privacy, security and safety, and workforce empowerment, and assigns governance responsibility to the Chief Information Officer, who designates the City Information Security Officer and the City Digital Privacy Officer to oversee compliance.

   The Policy prohibits five uses outright, with no exception path: real-time or covert biometric identification; emotion analysis; fully automated decisions substantially affecting individual rights or safety without meaningful human oversight; social scoring; and cognitive behavioral manipulation.

   For any system falling outside those five prohibitions, the Handbook establishes a risk-based review. Systems are classified as Low, Medium, or High risk, based primarily on the affected individual's ability to opt out and the severity of potential harm. Low-risk systems do not require a full review. Medium-risk and High-risk systems require an AI Review coordinated by the City Digital Privacy Officer, including an Impact Assessment Form, a vendor-completed AI Fact Sheet, and, for higher-risk or high-public interest systems, a Data Usage Protocol and possible escalation to a Council Committee at the Chief Information Officer's discretion. Where a system is procured through a Request for Proposals, the Handbook further requires that AI-specific criteria be built into the RFP and scored using a ten-category framework covering system overview, training data, system operations, performance, ethical considerations, reliability, interpretability, monitoring, transparency, and user interaction.

2) **Individual Use of Generative AI (GenAI).** The City's Generative AI Guidelines govern staff use of tools such as ChatGPT, Claude, and Microsoft Copilot Chat for daily work. The baseline rule applies regardless of tool or account tier: absent Information Technology Department (ITD) approval and use of a City enterprise GenAI tool, no prompt may contain information not suitable for public release. Staff must use a City-specific account for City-related use, review and fact-check all output, and record at least one instance of each type of GenAI use through the City's Generative AI Usage Form.

   The Guidelines assess risk in two dimensions: the risk that
   - information exchanged with the tool reaches an unintended audience
   - the tool's output causes an adverse impact if relied upon for a decision, task, or service

   Two categories of use are prohibited unless the ITD approves an exception: using GenAI to evaluate or decide matters concerning individuals, including job applicants, proposals, or existing employees; and using GenAI to generate a person's voice or other synthetic audio.

ORS's Current Use of AI

Usage of GenAI across Administration, Executive Management, Investments, Accounting, Benefits, and Information Technology is informal, staff-directed, and limited to drafting, research, summarization, and quality review. Accounting has used GenAI to create “macro”- type functions to standardize data inputs for reporting. Each division reports that it requires human review of outputs before using them to support Board materials, communications, or decisions. Information Technology is also developing a tool to extract unstructured information from resolution PDFs to create a searchable archive.

The Investment Office is also onboarding EnvestHub, a workflow management system supporting pipeline management, due diligence, and portfolio monitoring. EnvestHub's current use case is limited to efficiency, extracting and categorizing information and producing draft summary materials for staff review, while the system is configured and staff are trained. Every stage requires staff verification, and an investment does not proceed without direction conversations, full team discussions, and Chief Investment Officer approval. Should the team later extend EnvestHub to rate or prioritize individual opportunities ahead of human review, that would be a materially different use case and would return through the internal review process described below before implementation.

Strengthening ORS's AI Practice

Measuring against the framework described above, staff identified four areas where current practice has not kept pace with the City's requirements.

First, ORS's vendor services contracts have not been systematically reviewed or updated to include the AI disclosure, audit rights, and data use protections the City's framework contemplates.

Second, the City's AI Handbook requires that AI-specific requirements be considered at the start of and during the procurement of AI-enabled solutions. While these were not incorporated into the most recent procurement from the outset, ORS worked with the City to complete its AI Risk Review for the system before contract execution.

Third, ORS staff has not regularly filed the City's Generative AI Usage Form. Because that reporting is self-administered and anonymous, ORS has no independent way to verify compliance directly. ORS will train employees on the requirement going forward. More significantly, the internal procedure proposed below does not depend on self-reporting, and will independently document Medium and High-risk AI use cases at ORS.

Fourth, ORS has not formally designated a departmental AI Lead, a role the Handbook establishes for each City department. This role interfaces with the City Digital Privacy Officer and raises department-level concerns and questions. In practice, the ORS IT Manager filled this role, but it was never officially designated.

Training Provided and Training Needs

Staff training to date addresses two distinct needs: general staff education in the safe and effective use of GenAI, and additional training specific to ORS's control over the issuance of paid GenAI licenses.

At ORS's June 2026 Quarterly All Staff meeting, "Using AI at Work" was presented to the entire ORS staff, covering the distinction between GenAI and traditional search, account setup and data privacy controls for GenAI tools, the categories of AI risk (inaccuracy and hallucination, misinformation, algorithmic bias, lack of transparency, and information leaks), the City's Generative AI Guidelines, and a structured framework for writing effective prompts.

Since August 2026, ORS has required, as a condition of approving a paid GenAI license, that the requestor complete self-paced training covering foundational use of the tool and two additional topics, typically safe and appropriate AI use and effective prompt construction, drawn from the City's approved list of artificial intelligence courses, together with a written business justification addressing why a paid license is necessary and why Microsoft 365 Copilot Chat does not meet the need. The City's Digital Privacy Office recently adopted the same standard as part of its transition from Enterprise ChatGPT to Enterprise Claude, now making it a city requirement to complete 3 training classes before provisioning a paid GenAI license. The Learning Management System tracks completion for all employees; ORS supervisors can view completion records for their direct reports.

Peer Retirement System Practice

Staff reviewed AI governance practices across a broad range of public retirement systems. Publicly available, adopted AI policies were difficult to identify within ORS's normal California peer groups. Staff therefore expanded the review and identified six public pension systems with publicly available AI policy or related documentation: SamCERA, SJCERA, Ohio SERS, OCERS, LACERA, and ACERA. The documents range from formally adopted policies to a public policy summary and an informational presentation.

The City's framework is more developed in several areas than the practices documented in these peer materials. In particular, the City's Generative AI Guidelines use a structured risk framework that evaluates both the risk of information breach and the risk of adverse impact, while the City AI Policy also requires an AI review, vendor disclosure of AI use, technical documentation, and an AI Fact Sheet when requested.

Staff also reviewed how the six systems assign authority to approve individual AI tools or use cases. Four of the six documents identify an operational approval authority. Ohio SERS assigns approval to a management and staff AI Oversight Committee, OCERS assigns approval to an AI Committee, and SJCERA assigns approval to Executive Leadership. SamCERA assigns CEO approval where Confidential Information will be used in a closed, secure AI system. Appendix A summarizes this comparison in full.

The peer documents reviewed, the City's AI framework, and ORS's proposed AI procedures do not define specific records retention requirements for AI prompts, outputs, and information retained within AI platforms. These references do not establish AI-specific retention categories or retention periods. The City's Generative AI Guidelines similarly state that retained City work must comply with existing City retention policies without defining how those policies apply to individual prompts or outputs. Staff will work with legal counsel to develop appropriate guidance as relevant fact patterns are identified.

Proposed Internal AI Governance Procedures at ORS

To close the gaps identified above and to govern ORS's use of AI going forward, staff will establish the following procedures:

1. ORS designated the IT Manager as the AI Lead for ORS.
2. ORS will establish an AI Oversight Committee (“AOC”), composed of the IT Manager, the Deputy Director, and the Deputy Chief Investment Officer.
3. The current department AI Use Case inventory will be verified by the AOC, and risk categorizations will be assigned.
4. ORS will review vendor contracts at each contract's renewal or optional extension, incorporating the AI disclosure, audit rights, and data use protections the City's framework contemplates.
5. Each new AI use case, whether a GenAI use case or an AI-enabled system procurement, will be submitted to the AOC for review, risk level confirmation, and recommendation to the CEO.
6. The AI Lead will refer appropriate use cases to the City's own AI Review process.
7. Any new Medium-risk use cases, and all High-risk use cases will be presented by the AOC to the Boards annually, timed to align with the City's own annual reporting cycle for AI systems requiring annual review.
8. Lastly, ORS will continue to target training to all staff around the AI procurement and contracting requirements, ORS Use Case approval process, AI Citation requirements, and GenAI safe use standards.

CONCLUSION

The City of San José's AI governance framework is sound and complete, addressing both the procurement and deployment of AI-enabled systems and the individual use of GenAI tools, with clear guiding principles, defined prohibitions, and a risk-based review process. Staff has measured ORS's current practice against that framework, identified the gaps described above, and is implementing the procedures described above to close them and govern ORS's use of AI going forward.

Staff recommends that the Board adopt City of San José Artificial Intelligence Policy 1.7.12 as ORS's governance baseline in this area. Formal adoption of Policy 1.7.12 is being brought forward as a separate action item under New Business at this meeting.

John Flynn
Chief Executive Officer

Attachments

1. City of San José Artificial Intelligence Policy 1.7.12
2. City of San José AI Handbook
3. City of San José Generative AI Guidelines

APPENDIX A: ORS COMPILED DATA

AI Use Case Readout

Use Case	Description	Functional Areas Involved	Current Tools	Data Involved and Safeguards	Risk Tier
Correspondence and Document Drafting	Drafting and editing emails, memoranda, guidelines, procedures, and public-facing documents	Administration, Executive Management, Information Technology, Accounting	Free or Paid GenAI	Public information and Board materials; PII and confidential information redacted; staff review for accuracy and tone before use	Lower
Document and Contract Review	Reviewing agendas and minutes for consistency, comparing document versions, comparing contract language against certificates of insurance, formatting for accessibility compliance	Administration, Accounting	Free or Paid GenAI	Public information, contracts, policies; staff verify against source documents before use in Board materials	Lower
Research and Summarization	Governance practice research, tax topic research verified against IRS or industry sources, legal terminology translation, summarizing lengthy policies or memoranda	Administration, Accounting, Executive Management, Information Technology	Paid GenAI	Public information; findings verified against official sources before reliance	Lower
Coding and Reporting Automation	Generating Power Query M code to automate reporting processes, advanced Excel formula support, development of the Resolution PDF Extractor	Accounting, Information Technology	Paid GenAI	Financial data from FMS, anonymized with no PII; public resolution documents; Accounting applies its existing preparer, senior accountant, and manager review structure	Moderate
Benefits and Member Support Research	Basic research, analysis, and document review supporting member communications and benefit administration	Benefits Operations	Free or Paid GenAI	Internal operational information and anonymized member information; output reviewed and edited before use	Moderate
Embedded Investment Workflow System	Pipeline management, due diligence, checklist enforcement and approval, contracting, portfolio monitoring, and operations, including AI assisted information extraction, first draft investment memoranda, and natural language queries	Investment Office	EnvestHub	Due diligence materials, financial statements, third party data, executed contracts, notes and transcripts; staged workflow requires staff verification and approval before advancing	Higher

Peer Research Summary

Topic	SamCERA	OCERS	Ohio SERS	SJCERA	LACERA
Document status	Adopted Board policy, April 28, 2026	Employee handbook provision, presented December 3, 2025	Adopted policy, effective November 29, 2023	Adopted Board policy, December 12, 2025	Public summary of a policy; internal Approval Level listed as Chief Executive Officer, last updated January 22, 2026
Board approval of individual AI tools or use cases required?	No; CEO approves closed, secure systems using Confidential Information	No; AI Committee reviews and approves use cases	No; AI Oversight Committee approves	No; Executive Leadership approves	Not addressed; internal procedures not described in the public summary
Board's role distinguished from individual use case approval?	Yes; Board of Retirement policy with broad fiduciary responsibility; CEO approves the specified Confidential Information use	Not addressed; policy does not assign either role to the Board	Not addressed in the material reviewed	Yes; Board adopted the policy and must approve amendments; Executive Leadership approves use cases	Not addressed; policy Approval Level is Chief Executive Officer, no Board role stated
Higher-risk escalation path	Terminates with CEO approval in consultation with staff for Confidential Information; no formal risk tier defined	Not addressed; no formal higher-risk tier defined	Routed through the AI Oversight Committee; no Board escalation tier	Not addressed; no higher-risk tier defined	Not addressed in the public summary
Ongoing Board involvement with individual use cases	Periodic educational sessions for the Board required; no individual use case approval or reporting	Not addressed; no periodic Board reporting stated	Not addressed in the material reviewed	Not addressed; no periodic Board reporting or approval stated	Not addressed in the public summary
Named approval roles; trustee participation	CEO named for Confidential Information approval; no trustee named in an operational approval role	AI Committee named; membership and any trustee role not disclosed	AI Oversight Committee, including the Executive Director/Deputy Executive Director, General Counsel, Chief Risk Officer, CTO, and CFO; no trustee named	Executive Leadership named; individual titles and any trustee role not disclosed	Information Security Officer named as Authorizing Manager; Chief Executive Officer named as Approval Level; no trustee named

ACERA's reviewed material is an operational presentation rather than an adopted policy. It confirms that ACERA has an AI Committee, an AI Policy, AI Guidelines, and AI language incorporated into contracts, but it does not disclose the Committee's composition or an individual use case approval process. The New York State Common Retirement Fund's 2026 Stewardship Priorities document (April 2026) addresses AI governance expectations the Fund applies to its portfolio companies through its stewardship and proxy voting program, not an internal policy for the Fund's own AI use, and is noted here for completeness rather than as a direct comparison point.

Source Document List

- City of San José Artificial Intelligence Policy 1.7.12, effective June 28, 2024. The City's citywide administrative policy establishing guiding principles, an AI Review requirement, vendor disclosure obligations, prohibited uses, and incident reporting.
- City of San José AI Handbook dated March 18, 2024. The implementing handbook beneath Policy 1.7.12, detailing the AI Review process, Low/Medium/High risk tiers, and the RFP evaluation framework for AI-enabled system procurements.
- City of San José Generative AI Guidelines, updated April 24, 2025. The operational guidance document beneath Policy 1.7.12, covering account setup, citation practices, a two-factor risk assessment framework, and risk-classified use case examples.
- City of San José Data Governance Policy 1.7.13, effective June 3, 2025. The companion citywide policy establishing data governance roles and departmental structures, adopted separately by the Federated Board on June 26, 2025, and the Police and Fire Board on August 7, 2025.
- SamCERA Policy for the Use and Deployment of Artificial Intelligence Technologies, adopted April 28, 2026. An adopted Board policy with express fiduciary framing and closed system requirements for confidential information.
- OCERS Artificial Intelligence Policy, presented to the Personnel Committee December 3, 2025. An Artificial Intelligence Policy included within the OCERS Employee Policy Handbook.
- Ohio SERS AI Usage Policy, effective November 29, 2023. An adopted policy establishing an AI Oversight Committee and a formal use case intake and approval process.
- SJCERA Artificial Intelligence Policy, adopted December 12, 2025. An adopted Board policy covering employees, contractors, and third-party service providers.
- LACERA Artificial Intelligence Policy Public Summary, last updated January 22, 2026. A public summary listing the Chief Executive Officer as the policy's internal Approval Level; the summary states that it does not describe internal procedures.
- ACERA, AI at ACERA, presentation dated December 18, 2025. An operational presentation describing current AI practices rather than an adopted policy.
- New York State Common Retirement Fund, 2026 Stewardship Priorities, April 2026. An investment stewardship document addressing AI and cybersecurity expectations for portfolio companies, not an internal AI use policy.

APPENDIX B: REFERENCE

Glossary of Key Terms

Term	Definition	Source
Artificial Intelligence (AI)	A machine-based system that, for a given set of human-defined objectives, can make predictions, recommendations, or decisions influencing real or virtual environments	Policy 1.7.12, Terms and Definitions
Algorithm	A series of logical steps through which an agent, typically a computer or software program, turns inputs into outputs	Policy 1.7.12, Terms and Definitions
System	Any software, sensor, or process that uses AI to automatically generate outputs, including predictions, recommendations, or decisions that augment or replace human decision-making	Policy 1.7.12, Terms and Definitions
AI Fact Sheet	A template capturing the essential technical details of an AI system, completed by vendors during procurement and maintained by ITD	Policy 1.7.12, Terms and Definitions
Generative AI (GenAI)	A machine that automatically creates content such as text, audio, or images	Guidelines, Definitions
Machine Learning	A type of AI in which computers use data to learn tasks through algorithms	Guidelines, Definitions
Risk of Information Breach	The potential harm if information exchanged with GenAI is released to an unintended audience, including personally identifiable information, sensitive records, or confidential business information	Guidelines, Assessing Risk in GenAI Use Cases
Risk of Adverse Impact	The potential harm of using GenAI output for a decision, task, or service, considered from an equity lens across different populations	Guidelines, Assessing Risk in GenAI Use Cases