I. Overview

California is home to global AI innovation, yet it also leads in regulatory experimentation with landmark bills like SB-1047 and SB-53. These laws establish a rigorous compliance regime for developers of powerful frontier models, including mandates for shutdown mechanisms, third-party audits, and whistle-blower protections. While crucial for public safety, these measures may unintentionally sideline academic researchers and startups who cannot meet such thresholds.

To reconcile innovation with public safety, we propose the California AI Sandbox: a public-private test-bed where vetted developers can trial AI systems under tailored regulatory conditions. The sandbox will be aligned with CalCompute infrastructure and enforce key safeguards including third-party audits and SB-53-inspired whistle-blower protection to ensure that flexibility does not come at the cost of safety or ethics.

II. Policy Recommendation

Establish the California AI Sandbox under the Government Operations Agency, integrated with CalCompute.

Key features:
- Eligibility: Startups, universities, and developers working on sub-threshold AI models (i.e., below SB-1047's $100M compute cost definition of "covered models").
- Access to CalCompute: Selected teams receive compute access to reduce infrastructure barriers.
- Time-bound participation: Up to 12 months of sandbox participation with a clearly defined scope, goals, and exit criteria.
- Third-party audits: Mandatory pre- and post-sandbox risk audits based on SB-1047 audit frameworks.
- SB-53 compliance: All participants must adopt internal whistleblower policies and protections mirroring SB-53, including anonymous reporting channels.
- Public reporting: Aggregate insights and safety learnings will be shared in annual reports to inform future AI regulations.

III. Evidence and Rationale

- Barriers for small innovators: SB-1047's thresholds (>10^{26} FLOPs, >$100M compute cost) are unreachable for most academic and early-stage teams. A sandbox allows regulatory breathing room without weakening oversight.
- Precedent for sandboxes: The UK, Singapore, and Canada have piloted AI and fintech sandboxes that enabled safe experimentation while informing regulatory policy.

- Complementary to CalCompute: SB-1047 and SB-53 already envision CalCompute as an equitable public resource. The sandbox proposal operationalizes this vision.
- SB-53 as a model: Whistleblower protections in SB-53 provide a tested structure for internal risk reporting. Embedding these policies in sandbox projects protects workers and strengthens accountability.

IV. Projected Impact

Area	**Expected Outcome**
Innovation	Enables responsible AI development by smaller actors
Policy evolution	Informs real-world regulatory design through sandbox learning
Equity	Expands access to public compute and research resources
Safety	Mandates audits and risk disclosures even under relaxed rules

V. Challenges and Mitigation

Challenge	Mitigation Strategy
Unsafe models slipping through	Mandatory third-party audits and sandbox approval process
Industry capture or misuse	Public interest groups and UC reps on review panel
Whistle-blower vulnerability	Enforce SB-53 protections; require anonymous reporting lines
Funding for oversight	Use existing Cal-compute funding streams and state partnerships

Conclusion
By launching a tightly governed California AI Sandbox, the state can protect its leadership in ethical AI while ensuring smaller innovators aren't left behind. The sandbox complements existing legislation (SB-1047 and SB-53), bridges gaps between risk and progress, and models a proactive, bipartisan approach to AI governance.