NEW YORK'S RAISE ACT

A CRITICAL ASSESSMENT

by **John Eden** | Counsel

INTRODUCTION

On December 19, 2025, Governor Kathy Hochul signed New York's ***Responsible AI Safety and Education Act*** (the "RAISE Act" or "Act"), which imposes transparency, safety protocols, and incident reporting requirements on developers of large frontier artificial intelligence models. The Governor's office characterized the legislation as **nation-leading**. The Act is expected to take effect on January 1, 2027.

While the RAISE Act represents a serious effort to address the risks posed by advanced AI systems, a close reading of the legislation reveals structural problems that may undermine its effectiveness and create unintended consequences for both large incumbents and emerging companies. This alert summarizes the Act's key provisions and offers a critical assessment of its practical limitations.

SCOPE AND APPLICABILITY

The RAISE Act applies to large developers of frontier AI models that are developed, deployed, or operating in New York.

"Large developers" are defined as persons with more than $500 million in annual revenue. This threshold aligns with California's Transparency in Frontier Artificial Intelligence Act, enacted in September 2025. Colleges and universities engaged in academic research are excluded.

"Frontier models" are defined as AI models trained using more than 1026 floating point operations, or models trained from such systems through "knowledge distillation," provided that the compute costs of that technique exceed $5 million.

THE EXTRATERRITORIALITY PROBLEM

The Act's jurisdictional language, which covers technologies developed, deployed or operating in New York, is broad enough to theoretically capture foreign hyperscalers, including Chinese cloud and AI companies such as Alibaba Cloud, ByteDance, and Baidu. Yet the Act in its current form offers no mechanism for extraterritorial enforcement. If a Chinese frontier model is accessible to New York users but the developer has no U.S. presence, New York's Attorney General has no practical means of compelling compliance, obtaining records, or enforcing penalties.

This raises an obvious question: Does the RAISE Act create an asymmetric compliance burden that disadvantages domestic developers while leaving foreign competitors, including those operated by geopolitical adversaries, effectively untouched? If so, the Act may function less as a safety measure and more as a competitive drag on U.S. AI development.

THE KNOWLEDGE DISTILLATION TRAP

The Act's treatment of "knowledge distillation" models warrants particular scrutiny. Knowledge distillation is a technique where a larger AI model or its outputs is used to train a smaller model with similar capabilities. Under the RAISE Act, any model created through knowledge distillation with compute costs exceeding $5 million qualifies as a frontier model subject to the full compliance regime.

This threshold is low enough to capture well-funded startups that are nowhere near the $500 million revenue mark. Consider a startup with $30 million in capital that has spent $5 million on distillation-based training. That company would be required to develop and maintain written safety and security protocols addressing scenarios involving "the death or serious injury of at least 100 people or at least $1 billion in damages", which is the Act's definition of "critical harm".

The mismatch is stark. The Act requires a robust compliance infrastructure: detailed written protocols, annual reviews, five-year retention, and robust incident reporting systems. This compliance infrastructure is designed for enterprise-scale operations. Imposing these obligations on early-stage companies building capable but resource-constrained products may be disproportionate to the actual risk profile of those products, and could deter innovation at precisely the stage where experimentation is most valuable.

SAFETY PROTOCOLS AND TESTING REQUIREMENTS

The RAISE Act requires large developers to develop, publish, and maintain written safety and security protocols before deploying a frontier model. These protocols must address, among other things,

protections and procedures that reduce the risk of critical harm, cybersecurity protections against unauthorized access or misuse, and testing procedures to evaluate whether the model poses an unreasonable risk of critical harm.

Developers must document and retain detailed test results and implement safeguards to prevent unreasonable risk of critical harm.

THE SANDBOX PROBLEM

These testing requirements raise a practical question that the Act does not answer: How does one conduct meaningful sandbox-based testing for harms that, by definition, only manifest at scale in the real world?

The critical harm threshold set forth in the RAISE Act — 100 deaths or serious injuries, or $1 billion in damages — describes catastrophic outcomes that emerge from complex interactions between AI systems, users, and real-world conditions. A controlled testing environment cannot replicate the conditions under which such harms occur. Red-teaming exercises can identify certain failure modes, but they cannot simulate the emergent behaviors of a model deployed to millions of users across diverse and unpredictable contexts.

The result may be compliance theater: Developers will document testing procedures and retain records, but the testing itself may have limited predictive value for the catastrophic scenarios the Act purports to address.

INCIDENT REPORTING

The RAISE Act requires large developers to report a "safety incident" to the New York Attorney General and the Division of Homeland Security and Emergency Services within 72 hours of discovery. Reportable incidents include unauthorized access, model misuse, and critical control failures. Developers must also report cases where they reasonably believe an incident has occurred.

This 72-hour window is notably stricter than California's TFAIA, which allows 15 days and requires actual knowledge of an incident rather than reasonable belief.

THE TIMING TRAP

The 72-hour reporting requirement suffers from the same structural problem as most incident response laws: it forces companies to communicate with regulators before they fully understand what happened.

Complex security incidents, particularly those involving sophisticated AI systems, often take days

or weeks to investigate properly. Root cause analysis, log review, and impact assessment are time-consuming processes. A 72-hour deadline creates pressure to file preliminary reports based on incomplete information, which may later prove inaccurate or misleading.

This creates a no-win situation. File early with incomplete information, and risk being accused of providing inaccurate reports. Wait until the picture is clear, and risk penalties for missing the deadline. The Act does not appear to provide a safe harbor for good-faith preliminary reports that are later corrected, which compounds the problem.

Governor Hochul's signing memo acknowledged this tension, noting that developers "may describe limitations on their knowledge" of incidents involving models modified by third parties. But this does not resolve the fundamental timing problem for incidents that are simply difficult to diagnose quickly.

ENFORCEMENT AND PENALTIES

The New York Attorney General has exclusive authority to enforce the RAISE Act. Civil penalties are capped at $1 million for a first violation and $3 million for subsequent violations. Courts may also issue injunctive or declaratory relief.

A PENALTY STRUCTURE THAT SATISFIES NO ONE

The Act's penalty structure raises questions from both ends of the spectrum. For large developers (i.e., hyperscalers with revenues exceeding $500 million), a $1 million fine is immaterial. These are companies with annual revenues in the tens of billions of dollars. A million-dollar penalty is a rounding error, not a deterrent. If the goal is to influence the behavior of the largest AI developers, this penalty structure is unlikely to accomplish it.

For smaller companies, including startups that may be swept into the regime through the knowledge distillation provision, a $1 million penalty could be an existential threat. A company with $50 million in funding that faces a $1 million fine, or the threat of one, may be forced to abandon a product line or curtail development, even if the underlying conduct was a good-faith compliance failure rather than willful misconduct. Moreover, being fined early on in a startup's lifecycle could make fundraising more difficult down the track.

The result is a penalty regime that is too weak to deter the largest players and too harsh for smaller ones.

A DEEPER PROBLEM: RULES VS. STANDARDS

Beyond the penalty amounts, the RAISE Act reflects a broader choice in regulatory design that deserves scrutiny. The Act adopts a principles-based approach: developers must implement reasonable protections, conduct testing to identify unreasonable risks, and report incidents that could lead to critical harm. These standards are inherently flexible, which means their content will be determined over time through regulatory guidance and enforcement actions.

This approach places substantial discretion in the hands of the Attorney General's office and the new oversight body within the Department of Financial Services. Whether a given safety protocol is reasonable, or whether a particular model poses an unreasonable risk, will depend on judgments made by regulators after the fact.

An alternative approach, one that the New York State Legislature ("Legislature") did not adopt, would be to identify specific applications of AI that are categorically prohibited. Rather than asking whether a developer's safety protocols are reasonable, a rules-based regime would ask whether the developer is building something that appears on a schedule of prohibited applications. This approach has its own limitations and flaws, including the difficulty of anticipating harmful applications in advance. But it offers greater clarity to developers and reduces the risk of arbitrary or inconsistent enforcement.

The RAISE Act's choice of a standards-based regime is not inherently wrong, but companies subject to the Act should understand that compliance will require ongoing engagement with regulatory guidance as it develops, and that the line between compliant and non-compliant conduct will certainly shift over time.

FEDERAL PREEMPTION RISK

The RAISE Act may face a federal challenge. On December 11, 2025, President Trump issued an executive order announcing support for a minimally burdensome federal AI regulatory framework and directing the Department of Justice to challenge state laws deemed inconsistent with that policy.

The executive order frames state AI laws as creating a patchwork that could stifle innovation. Whether the Department of Justice will bring suit against New York and California remains to be seen, but companies should be aware that the RAISE Act's long-term viability is uncertain.

A CLOSING OBSERVATION: WHAT IS THIS LAW ACTUALLY FOR?

Practitioners and compliance officers will dutifully parse the RAISE Act's requirements and build

programs to satisfy them. But it is worth stepping back to ask a more fundamental question: What does the Legislature believe this law will accomplish?

If the goal is to prevent catastrophic harms from frontier AI systems, the Act's mechanisms seem poorly suited to the task. Testing requirements that cannot replicate real-world conditions, incident reporting timelines that incentivize incomplete disclosures, and penalties that are immaterial to the largest developers do not add up to an effective safety regime.

If the goal is to assert regulatory jurisdiction over a transformative technology, the Act is more fit for purpose. Having said that, having jurisdiction over a wide range of frontier AI companies without effective enforcement tools for US and non-US based companies could easily disadvantage U.S. entrepreneurs vis a vis their foreign counterparts. In short, there is serious cause for concern that the Act will effectively burden domestic developers while leaving foreign competitors untouched. If this concern materializes once the Act becomes law, the U.S. AI ecosystem could be, unfortunately, hobbled by design.

None of this means companies should ignore the RAISE Act. It is law, and compliance is required. But companies building frontier AI systems should approach this regime with clear eyes about its limitations—and should engage with policymakers as the regulatory landscape continues to evolve.

This altert is provided for informational purposes only and does not constitute legal advice. For questions about compliance with the RAISE Act or other AI regulatory matters, please contact ***John Eden*** *at jeden@burghergray.com. BurgherGray is a dynamic corporate boutique law firm comprised of highly experienced and diverse attorneys, most of whom have honed their skills practicing at large corporate law firms and in-house legal departments of large corporations and governmental agencies. The firm counsels clients ranging from emerging companies to Fortune 100 enterprises on a range of complex business litigation, government and internal investigations and transactional matters, including finance, M&A, securities regulation and corporate and commercial transactions. The firm is a member of NAMWOLF, a member of NMSDC, and is certified as an MBE by the City and State of New York and the City of Chicago and State of Illinois.*

BURGHERGRAY LLP
ATTORNEYS AT LAW

www.burghergray.com

1350 Broadway | Suite 1510
New York, NY 10018
T: 646.513.3231 | F: 646.561.9866
info@burghergray.com