Section-by-Section Summary

FRONTIER Act 2.0 (discussion draft)

Sec. 1. Short title.

This section provides that the Act may be cited as the Frontier Risk Oversight, National Transparency, Independent Evaluation, and Reporting Act of 2027 — the **FRONTIER Act of 2027**.

Sec. 2. Definitions.

This section defines the key terms used throughout the Act, including the following:

Frontier model / frontier developer — The term "frontier model" means a foundation model trained using more than 10^{25} operations, counting the original training run and any subsequent modification. The term "frontier developer" means any entity operating in or affecting interstate or foreign commerce that has trained, or initiated training of, such a model. Once a developer is covered, its obligations reach all of its foundation models, except that model cards remain limited to frontier models.

Large / very large frontier developer — These terms establish the higher coverage tiers, keyed to AI-related development expenditures and measured with affiliates over the preceding 36 months: at least $1B in such expenditures for the large tier, and at least $5B for the very large tier. Each higher tier carries every lower tier's duties.

Catastrophic risk — The term "catastrophic risk" means a foreseeable, material risk that a frontier developer's foundation model materially contributes to 50+ deaths or serious injuries, or $1B+ in property damage, in a single incident involving the model (i) providing expert-level assistance in developing or releasing a CBRN weapon or in conducting offensive cyber operations, (ii) engaging in conduct without meaningful human oversight that is a cyberattack or would constitute murder, assault, extortion, or theft, (iii) evading the control of its developer or user, or (iv) automating AI research and development in a way that substantially accelerates any of those risks. Because automated AI research and development, including recursive self-improvement (RSI), is its own category of catastrophic risk, each of the Act's obligations applies to it: frameworks must address it, large developers report quarterly on AI R&D capabilities and internal use, and very large developers must retain an IVO dedicated to assessing it.

Imminent catastrophic risk — The term "imminent catastrophic risk" means a present or impending catastrophic risk, and serves as the predicate for emergency orders and mandatory IVO referral.

Acceptable levels of catastrophic risk mitigation — The term means risk mitigation adequate to ensure that a foundation model's anticipated benefits outweigh its catastrophic risk, taking into account the probability and magnitude of each.

Critical safety incident — The term "critical safety incident" means unauthorized exfiltration or malicious modification of model weights; a cyberattack on the developer; harm from a materialized catastrophic risk; loss of control, including the model deceiving its developer to

subvert controls, outside a designed evaluation; a near miss; developer-identified incidents bearing on catastrophic risk or public safety; incidents affecting critical infrastructure; or incidents the Secretary designates by rule.

Near miss — The term "near miss" means an event in which a material safeguard, control, or monitoring system applied to a foundation model, including during testing or internal use, fails, is bypassed, or is materially degraded, and, but for timely intervention or luck, would likely have led to the compromise of model weights, harm from a materialized catastrophic risk, or loss of control, but no such harm resulted.

Secretary — The term refers to the Secretary of Commerce, acting through a new Office of AI Security (Sec. 3), who holds the Act's rulemaking authority, receives incident reports and registrations, and licenses and oversees IVOs.

Sec. 3. Office of AI Security.

This section establishes an Office of AI Security within the Department of Commerce, headed by a Senate-confirmed Assistant Secretary of Commerce for AI Security, through which the Secretary carries out the Act, provides hiring authority, and authorizes $100 million a year for the Office.

Sec. 4. Rulemaking and guidance.

This section vests the Secretary with enumerated notice-and-comment rulemaking authorities to carry out the Act. The Secretary must issue rules governing IVO licensing and oversight, including time-limited, category-specific licenses, graduated sanctions, and scaled fees, and setting criteria for substantial (model) and material (framework) modifications. The Secretary may also issue rules adjusting the compute and expenditure coverage thresholds and adding capability-based criteria for frontier models, calibrated to reach the developers whose activities present catastrophic risk, and governing incident reporting, without narrowing the statutory definition of a critical safety incident. It also directs the Secretary to issue nonbinding guidance, after public comment, on approaches large frontier developers may use to meet the Act's mitigation standard, along with competence standards for interim assessors and a model access plan for IVOs, reviewed annually in light of corrective orders.

Sec. 5. Cumulative obligations.

This section provides that each higher tier complies with every requirement applicable to the tiers below it.

Sec. 6. Registration and disclosure.

This section requires large frontier developers to register with the Secretary, as a condition of developing, deploying, or operating any foundation model, by filing and annually renewing a disclosure statement identifying the developer, its significant owners, and its points of contact. It directs the Secretary to collect fees to offset administrative costs and maintain the public registry of filed developers.

Sec. 7. Model cards.

This section requires all frontier developers to publish, before or when deploying a new or substantially modified frontier model, a model card on its capabilities, intended uses, and restrictions. Large frontier developers must also include their catastrophic-risk assessments and results.

Sec. 8. Incident reporting.

This section requires reporting of critical safety incidents to the Secretary within 72 hours of when the developer knew or reasonably should have known of them (near misses and other developer-identified incidents, once identified), or within 24 hours where an incident poses an imminent threat, with follow-up reports on the most serious incidents. Large frontier developers must also report quarterly on catastrophic risk arising from their foundation models, including internal use and capabilities that could automate or substantially accelerate AI research and development; file a rapid report within 72 hours of a substantial increase in that risk; and certify in each quarterly report that they have reported every critical safety incident known to them. A governor may designate a State emergency management agency to receive incident information affecting the State under a confidentiality agreement.

Sec. 9. Frontier AI frameworks.

This section requires large frontier developers to write, implement, comply with, and publish a frontier AI framework — covering catastrophic-risk thresholds and assessment, model-weight cybersecurity, incident response, and deployment and internal-use decisions, and incorporating widely-accepted national, international, and industry approaches, best practices, frameworks, and methods — that is reasonably designed to achieve acceptable levels of catastrophic risk mitigation for each of their foundation models. It requires annual review of the framework and an annual independent compliance audit, with publication of the redacted audit report and its transmittal to the Secretary.

Sec. 10. Developer redactions and protection of submitted information.

This section permits redactions in published material only as necessary to protect trade secrets, risk-prevention mechanisms, cybersecurity, public safety, national security, or privacy, with a justification for each redaction and an unredacted copy to the Secretary, who may require an overbroad redaction to be narrowed. Information submitted to the government is exempt from FOIA and similar State and local disclosure laws.

Sec. 11. Independent verification.

This section directs the Secretary to license and oversee IVOs. It requires each very large frontier developer to retain, for each category of catastrophic risk, an IVO for ongoing assessment of whether its framework, governance, risk-monitoring, and mitigations achieve acceptable levels of catastrophic risk mitigation. IVO selection phases in by category: the developer first chooses its own assessor that meets industry competence standards and the Act's independence requirements; once licensed IVOs have capacity for a category, the Secretary assigns the IVO; and as the market matures, Secretary assignment is combined with developer choice. Assigned IVOs are paid from a fee-funded pool; developer-selected IVOs are paid by the developer. Access is governed by a single access plan for each developer, proposed by its lead IVO and approved by the Secretary, covering unredacted materials, model weights, internal-use models,

and personnel, with a resident IVO team on site; the developer may object only on limited grounds before an administrative law judge.

The IVO determines the scope and frequency of its assessments, reporting at least every 6 months and out of cycle when risk changes, and recommends corrective actions. If the developer's response is inadequate and it does not cure the deficiency, the Secretary may, after notice and a hearing, order the developer to take necessary actions, subject to judicial review; compliance with an order satisfies the framework standard for that deficiency. It requires an IVO that determines a model poses an imminent catastrophic risk to refer the matter to the Secretary within 72 hours for consideration of a Sec. 15 emergency order, and to notify the Secretary of unreported critical safety incidents.

It sets statutory independence standards for IVOs, provides for regular inspection of IVOs, and voids engagement terms that tie payment to results or penalize adverse findings. It permits IVOs to subcontract specialized portions of an assessment while remaining fully responsible, prohibits material misrepresentations in any report or opinion, and protects IVOs from suit over good-faith findings and reports, other than for bad faith or willful misconduct.

Sec. 12. IVO capacity-building grants.

This section authorizes the Secretary to award grants and cooperative agreements to build IVO capacity, with priority for categories of catastrophic risk that have too few licensed IVOs and for risks from automated AI research and development, but not to pay for assessing any particular developer.

Sec. 13. Enforcement.

This section authorizes the Attorney General to enforce the Act through civil penalties and injunctions. Violations carry a civil penalty of up to the greater of $1M per violation or a percentage of global annual revenue, with higher caps for repeat violations, and claims that a framework is not reasonably designed require prior notice and an opportunity to cure. It penalizes knowing IVO misrepresentations, prohibits materially false or misleading statements about catastrophic risk or compliance, and requires the Secretary to refer IVO-reported violations to the Attorney General or explain in writing why not. Violations of a Sec. 11 corrective order are enforceable on the same terms, and the Attorney General may not bring a framework-reasonableness claim over an IVO-identified deficiency while the corrective-action process is pending or an order addressing it is in effect.

Sec. 14. Near-miss penalties.

This section authorizes the Secretary to assess a civil penalty against a frontier developer for each near miss involving its models. The Secretary sets penalty amounts by rule so that each developer's expected penalties equal the expected catastrophic harm its models create beyond what it could pay (its assets after such harm, plus insurance). Each near miss bears the expected harm on the pathway it reveals, such as loss of control or weapons misuse, divided by the number of near misses that pathway is expected to produce, so that, for a given level of risk, rarer near misses carry larger penalties. If the developer failed to report a near miss as Sec. 8 requires, the penalty is multiplied by 10, unless the developer reasonably concluded in good faith that the event was not reportable. Near misses in designed evaluations carry no penalty. A

developer may contest a penalty in district court, which reviews it de novo. Penalties are in addition to those under Sec. 13 and apply only to near misses involving models deployed or used internally after the rule takes effect.

Sec. 15. Emergency orders.

This section authorizes the Secretary of Commerce to suspend or restrict a frontier developer's development, deployment, or internal use of a foundation model upon a written finding that it presents an imminent catastrophic risk. A provisional order may issue on a preliminary determination — with notice and an opportunity to cure unless the imminence of the risk forecloses them — and expires after 45 days; a final order requires the written finding and may be renewed only on a fresh finding. Each order must state its basis, why a narrower measure would not suffice, the restrictions imposed, and the conditions for rescission, and the Secretary must rescind the order once those conditions are met or the risk has passed. The developer may seek an expedited administrative hearing at which the Secretary bears the burden of justifying the order, and expedited judicial review in the D.D.C., with appeal to the D.C. Circuit; to keep a longer order in place, the Secretary must prove its necessity in court. Congress receives each order and semiannual reports on the use of the authority. Violations carry a civil penalty of up to $10M per violation per day, and willful violations are criminal.

Sec. 16. Reports to Congress.

This section directs the Secretary to report annually to Congress on critical safety incidents and IVO-reported violations, and directs the Comptroller General to report annually to House Energy & Commerce and Senate Commerce on the state of the IVO market, IVOs' independence from the AI industry, and recommendations to preserve that independence.

Sec. 17. Relationship to State law.

This section provides that nothing in the Act preempts any State law at least as protective of persons or property as the Act, except for the confidentiality of federal materials shared with States and the protection for IVOs' good-faith performance.

Sec. 18. General provisions.

This section provides that the Act takes effect on enactment and includes a severability clause.

Sec. 19. Sunset.

This section provides that the Act ceases to have effect 10 years after enactment.