Harvard Journal of Law & Technology
Volume 39, Number 1 Fall 2025

COMPUTE MONITORING: REVIVING THE ENCRYPTION BACKDOOR DEBATE

*Lisa Lu**

ABSTRACT

As AI regulation advances globally, compute governance — the oversight of computational resources used to train large AI models — has emerged as a key regulatory strategy. This approach promises visibility into potentially harmful AI development but raises significant constitutional concerns when implementation requires access to encrypted user data or mandates the construction of decryption capabilities.

This Note examines the civil liberties implications of compute monitoring regimes under both Fourth and First Amendment jurisprudence. It argues that government-compelled access to user data through compute providers (like Amazon, Microsoft, or Google) may constitute an unconstitutional search, particularly where users have employed encryption and confidential computing to secure their information. The analysis shows that users maintain reasonable expectations of privacy even when using third-party computing services, when those users have taken affirmative steps to protect their data through technical safeguards.

Furthermore, this Note contends that requiring providers to build decryption mechanisms or backdoors constitutes compelled speech in violation of the First Amendment. By analyzing code as expressive conduct deserving constitutional protection, this Note argues that mandating the creation of surveillance infrastructure forces providers to "speak" in ways that may conflict with their values and promises to users.

The Note distinguishes between constitutionally permissible reporting of objective compute metrics — such as processing power utilized or training duration — and problematic demands for subjective assessments or invasive data disclosures that implicate protected rights. It concludes by urging policymakers to pursue narrowly tailored, privacy-preserving implementations of compute governance that achieve legitimate regulatory objectives without undermining fundamental constitutional protections.

---

* Stanford Law School, Class of 2025. I am grateful to Professor Daniel E. Ho for directing this research and for his guidance and feedback on this piece. I also thank the editors of the *Harvard Journal of Law & Technology* (particularly Ebun Ajayi) for their thoughtful comments and feedback. All errors are my own.

TABLE OF CONTENTS

I. INTRODUCTION

In recent years, decision-makers around the world have begun to consider and implement regulations for the development and usage of artificial intelligence ("AI") systems. This interest in AI regulation has largely been driven by the rapid emergence of foundation models — large-scale AI models that can be adapted to a wide range of downstream tasks — and by concerns over their potential misuse.[1] Governments have proposed and adopted a variety of regulatory approaches to address the perceived risks posed by these models.[2]

One such regulatory approach is compute governance — the use of controls over computational resources as a way to regulate the development and deployment of advanced AI systems.[3] This approach is appealing because compute is measurable, highly concentrated in the AI supply chain, and directly tied to the infrastructure required to train and deploy large AI models.[4] By monitoring how large amounts of computing resources are being used, policymakers can thus gain visibility into potentially problematic AI usages and respond more quickly.[5]

However, that heightened visibility could prove too intrusive and risk infringing on civil liberties. The top compute providers go to great lengths to preserve the privacy of their customers' data,[6] offering security measures like encryption and accompanying assurances in their policies explicitly designed to protect user confidentiality.[7] But governments could require compute providers to access and disclose data beyond mere compute usage or transactional information, for the stated purpose of ascertaining whether the compute is being used to develop a model with dangerous capabilities or a similarly-sized, but harmless,

---

1. Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx et al., *On the Opportunities and Risks of Foundation Models* 3 (July 12, 2022) (unpublished manuscript) (on file with arXiv), https://arxiv.org/pdf/2108.07258 [https://perma.cc/888H-JJRD].

2 . Rishi Bommasani, Sayash Kapoor, Kevin Klyman, Shayne Longpre, Ashwin Ramaswami, Daniel Zhang et al., *Considerations for Governing Open Foundation Models*, 386 SCI. 151, 151 (2024).

3 . Lennart Heim, Markus Anderljung, Emma Bluemke & Robert Trager, *Computing Power and the Governance of AI*, CTR. FOR GOVERNANCE AI (Feb. 14, 2024), https://www.governance.ai/analysis/computing-power-and-the-governance-of-ai [https://perma.cc/T5NQ-Z3DT]. "Compute" refers to the amount of processing power needed to train or run an AI model, measured in the number of computations a given resource can execute in a second. Jai Vipra & Sarah Meyers West, *Computational Power and AI*, AI NOW INST. (Sep. 27, 2023), https://ainowinstitute.org/publications/compute-and-ai#h-what-is-compute-and-why-does-it-matter [https://perma.cc/F88E-GHX5]. The term is generally inclusive of both hardware resources (e.g., chips) and the supporting software required for the particular task. *Id.*

4. Heim et al., *supra* note 3.

5. *Id.*

6. *Id.*

7. *See* discussion of privacy commitments *infra* Section III.A.

model.[8] Such a requirement could effectively force compute providers to create a government backdoor into data stored on their platforms, running afoul of the First Amendment's prohibition on compelled speech and challenging long-standing expectations of privacy in the cloud under the Fourth Amendment.[9]

This Note explores the legal claims that compute providers or the customers of compute providers might raise under the Fourth and First Amendments, respectively. Specifically, it assesses the strength of these claims based on the type of information that could be sought under a compute governance regime, considering both scenarios where access is limited to compute usage and transactional data (such as user identity and payment information), and those where access extends to underlying training data or other cloud-hosted content. It concludes by considering how a compute governance regulation could reconcile a government's interest in safety and customers' interests in privacy.

II. IMPLEMENTATIONS OF COMPUTE GOVERNANCE

As explained above, compute governance is a regulatory approach that may allow governments to track — and control — AI development with greater visibility. Compute has emerged as a way to classify AI systems by their potential to cause harm, as part of a tiered approach.[10] Based on how an AI system is classified, a regulated entity could be subject to greater reporting requirements and scrutiny.[11]

Compute monitoring provisions are being proposed and adopted worldwide.[12] In the United States, a compute governance approach to AI regulation received formal recognition when the Biden

8. Heim et al., *supra* note 3.

9. Diane Bernabei, James N. Baker & Cosimo L. Fabrizio, *Legal Challenges to Compute Governance*, LAWFARE (May 16, 2024, at 10:07 ET), https://www.lawfaremedia.org/article/legal-challenges-to-compute-governance [https://perma.cc/N3MH-62VK].

10. *See, e.g.*, *EU AI Act: First Regulation on Artificial Intelligence*, EUR. PARL. (Aug. 6, 2023), https://www.europarl.europa.eu/topics/en/article/20230601STO93804/eu-ai-act-first-regulation-on-artificial-intelligence [https://perma.cc/EJ92-3VYT] (describing how AI Act will establish differing obligations for providers and users depending on level of risk posed by AI system); Exec. Order No. 14,110, 3 C.F.R. 657, 665 (2024) (imposing reporting requirements for dual-use foundation models, which among other criteria could "pose a serious risk to security, national economic security, national public health or safety, or any combination of those matters"). Such tiering systems have their own policy and feasibility challenges, which are outside the scope of this piece and its focus on the legal challenges. *See* Rishi Bommasani, *Drawing Lines: Tiers for Foundation Models*, STAN. CTR. FOR RSCH. ON FOUND. MODELS (Nov. 18, 2023), https://crfm.stanford.edu/2023/11/18/tiers.html [https://perma.cc/936B-ZH2Y].

11. *EU AI Act: First Regulation on Artificial Intelligence*, *supra* note 10.

12. Bernabei et al., *supra* note 9.

Administration signed a landmark Executive Order on AI on October 30, 2023,[13] though that Order was one of many revoked by the following Trump Administration.[14] Among its many provisions, the Order tasked compute providers with alerting the government whenever foreign users utilize their products to train models that could be used for malicious activity, based on a compute threshold to be set by the Secretary of Commerce.[15] On January 29, 2024, the Department of Commerce's Bureau of Industry and Security issued a proposed rule implementing the compute monitoring provisions for IaaS providers in that Executive Order and a preceding Executive Order issued under the first Trump Administration.[16] Under the draft rule, domestic compute providers had to "report to the Department information on instances of training runs by foreign persons for large AI models with potential capabilities that could be used in malicious cyber-enabled activity."[17] The report would include the existence of the training run and identifying information, including "the customer's name, address, the means and source of payment for the customer's [a]ccount, email addresses, telephone numbers, and IP addresses."[18]

The draft rule further defined the scope of its reporting requirements as applying to any AI model (1) with the technical conditions of a "dual-use foundation model," or "technical parameters of concern," and (2) "capabilities that could be used to aid or automate aspects of malicious cyber-enabled activity."[19] Dual-use foundation models are general-application models trained on broad data and containing at least tens of billions of parameters.[20] Examples of malicious activity under this definition included social engineering attacks, denial-of-service attacks, and generation of misinformation.[21] The draft rule further left open how a model would be deemed potentially capable of being used for malicious purposes, leaving that authority to the Secretary.[22]

---

13. Rishi Bommasani, Christie Lawrence, Lindsey Gailmard, Caroline Meinhardt, Daniel Zhang & Peter Henderson et al., *Decoding the White House AI Executive Order's Achievements*, STAN. INST. FOR HUM.-CENTERED A.I. (Nov. 2, 2023), https://hai.stanford.edu/news/decoding-white-house-ai-executive-orders-achievements [https://perma.cc/2MZH-VDK2]; *see also* 3 C.F.R. 657 (2024).

14. Exec. Order No. 14,148, 90 Fed. Reg. 8237, 8240 (Jan. 20, 2025). The Trump Administration indicated it would be further assessing which components of the Biden Executive Order to retain. Exec. Order No. 14,179, 90 Fed. Reg. 20 (Jan. 23, 2025).

15. 3 C.F.R. 665 (2024).

16. Taking Additional Steps to Address the National Emergency with Respect to Significant Malicious Cyber-Enabled Activities, 89 Fed. Reg. 5698, 5706 (proposed Jan. 29, 2024) (to be codified at 15 C.F.R. pt. 7).

17. *Id.* at 5706. The proposed rule defines a "training run" as a person's transaction with a compute provider to train an AI model. *Id.* at 5702.

18. *Id.* at 5706.

19. *Id.* at 5702.

20. *Id.* at 5725.

21. *Id.* at 5702.

22. *Id.*

Compute monitoring provisions were also proposed at the state level. On February 4, 2024, California State Senator Scott Wiener introduced an AI safety bill that defined frontier AI models with a compute threshold.[23] "Covered models" exceeding the compute threshold were subject to requirements like additional due diligence, testing, and determination findings.[24] The bill was ultimately vetoed.[25]

Elsewhere in the world, the European Parliament and the European Council adopted the Artificial Intelligence Act on June 13, 2024.[26] The Act imposes strict requirements on models that are deemed to pose "systemic risk," similarly based on a compute threshold.[27] China may also be considering compute-based measures, as indicated by recent policy discussions.[28]

While many approaches to compute governance proposed thus far have primarily relied on compute-based thresholds, the U.S. approach — particularly as outlined in the Department of Commerce's draft rule — went further. The draft rule required compute providers to flag training runs involving large AI models based on the models' capabilities, without indicating how such capabilities would be determined and by who. Compliance with such a requirement thereby implicates not just compute usage but also the underlying model inputs, metadata, or training data. This broader reach raises concerns around users' privacy rights, compute providers' privacy guarantees, and the limits of government access to cloud-based data. The next section explores these implications through the lens of the Fourth and First Amendments.

III. FOURTH AND FIRST AMENDMENT CHALLENGES TO COMPUTE MONITORING

As compute monitoring provisions evolve from policy proposals into real-world implementations, they raise legal questions about the scope of government access to information stored by compute providers for their users. Though compute monitoring provisions are framed as enhancing public safety and national security, [29] their

23. S.B. 1047, 2023–2024 Reg. Sess. § 3 (Cal. 2024) (as introduced in Senate, Feb. 7, 2024).

24. *Id.*

25. Matteo Pistillo, Suzanne Van Arsdale, Lennart Heim & Christoph Winter, *The Role of Compute Thresholds for AI Governance*, 1 GEO. WASH. J.L. & TECH. 26, 29 (2025).

26. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 Laying Down Harmonized Rules on Artificial Intelligence and Amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act) (Text with EEA Relevance), art. 51, 2024 O.J. (L 1689) 1, 123.

27. *Id.*

28. *See* Pistillo, et al., *supra* note 25, at 29–30.

29. Exec. Order No. 14,110, 3 C.F.R. 657 (2024).

implementation may require compute providers to peer into sensitive customer data beyond compute usage. Such a requirement would conflict with the technical features and contractual guarantees offered by the predominant compute providers, which often advertise strong privacy protections, encryption, and customer data ownership.[30]

This Section analyzes legal challenges to potential compute monitoring provisions. It begins with an overview of the technical and contractual assurances set forth by the dominant compute providers that establish their users' expectations of privacy. This Section then turns to potential claims by users under the Fourth Amendment, focusing on how mandated reporting under a compute monitoring provision might run afoul of the Amendment's prohibition against unreasonable searches. Finally, it considers potential First Amendment claims from compute providers themselves.

A. Technical and Contractual Guarantees by Compute Providers

The dominant companies in the cloud market, Amazon, Microsoft, and Google, have all experienced increased growth in response to the excitement around generative AI.[31] These compute providers are alike in the technical and contractual guarantees they offer to their users in two ways. First, these providers all offer encryption and confidential computing services, allowing users to secure their data through the underlying technology. Second, these providers have established policies limiting the scenarios in which they may access customer data.

1. Technical Safeguards Rely on Encryption and Confidential Computing.

From a technical standpoint, all three compute providers offer encryption and confidential computing to their users. Also, the providers have all adopted a shared responsibility model in which they are responsible for the security of the cloud infrastructure, while the customers are responsible for securing their own applications and data.[32]

Microsoft Azure, Microsoft's cloud computing platform, encrypts data at rest and in transit.[33] Customers have the option to keep their own encryption keys, such that Microsoft does not have access to the keys

30. *See infra* Section III.A.

31. *Cloud Market Gets Its Mojo Back; AI Helps Push Q4 Increase in Cloud Spending to New Highs*, SYNERGY RSCH. GRP. (Feb. 1, 2024), https://www.srgresearch.com/articles/cloud-market-gets-its-mojo-back-q4-increase-in-cloud-spending-reaches-new-highs [https://perma.cc/F5H4-P78B].

32. Sina Ahmadi, *Systematic Literature Review on Cloud Computing Security: Threats and Mitigation Strategies*, 15 J. INFO. SEC. 148, 154 (2024).

33. *What is Confidential Computing?*, MICROSOFT (May 7, 2025), https://learn.microsoft.com/en-us/azure/confidential-computing/overview [https://perma.cc/TT2Q-PNMT].

or the encrypted data.[34] Users can further enable confidential computing. Confidential computing limits access to data in use, including from the cloud operator, by performing computation in a secure and isolated hardware-based environment.[35] When enabled, Microsoft cannot access unencrypted customer data either.[36]

Under Amazon's "Shared Responsibility Model," its customers are responsible for managing and encrypting their data.[37] Customers retain full control of content they upload to its cloud computing platform, Amazon Web Services ("AWS"), including how such content is stored and secured, and can also keep full control over their encryption keys.[38] Thus, customers can ensure that only they have access to their encrypted data.[39] AWS further offers a confidential computing system, which is designed to have no cloud operator access.[40] AWS operators cannot access any data stored on the instance storage, and if maintenance work is needed, the operator can do so only by using a strictly limited set of APIs.[41] These APIs also lack the ability to access customer data, and no AWS operator can bypass these protections.[42] The confidential computing system, therefore provides customer protection from the compute provider.

Google has a more involved "shared fate" model, which it distinguishes as ensuring customers can operate securely in the cloud by default, in addition to providing them the services to secure their data

---

34 . *About Our Practices and Your Data*, MICROSOFT, https://blogs.microsoft.com/datalaw/our-practices/#what-do-you-do-encryption-keys [https://perma.cc/5LU9-9L32].

35. MICROSOFT, *supra* note 33.

36. *Id.*

37 . Shared Responsibility Model, AMAZON WEB SERVS., https://aws.amazon.com/compliance/shared-responsibility-model/ [https://perma.cc/HN5W-Z76F].

38. Data Privacy FAQs, AMAZON WEB SERVS., https://aws.amazon.com/compliance/data-privacy-faq/ [https://perma.cc/ACA6-Y7VB].

39. Notably, AWS decided to encrypt new objects uploaded to S3, its cloud storage system, automatically beginning January 5, 2023, shifting from an opt-in model of encryption to an opt-out one. Sébastien Stormacq, *Amazon S3 Encrypts New Objects by Default*, AMAZON WEB SERVS. (Jan. 5, 2023), https://aws.amazon.com/blogs/aws/amazon-s3-encrypts-new-objects-by-default/ [https://perma.cc/ULT7-ATW8].

40. David Brown, *Confidential Computing: An AWS Perspective*, AMAZON WEB SERVS. (Aug. 24, 2021), https://aws.amazon.com/blogs/security/confidential-computing-an-aws-perspective/ [https://perma.cc/C5EZ-2DRY].

41. *Id.* An API or application programming interface is a mechanism that allows two software components to communicate with one another through requests and responses. *What is an API (Application Programming Interface)?*, AMAZON WEB SERVS., https://aws.amazon.com/what-is/api/ [https://perma.cc/ZC3X-6V9S]. The client application sends a request to the server, which returns a response. *Id.* Because APIs can be configured to verify the identity of the client sending a request, to determine what actions or data the client has access to, and to serve as a limited or even single point of entry to certain services or data, they can implement fine-grained access control and thereby support security and privacy. *See* Sandeep Kumar Jangam, Nagireddy Karri & Partha Sarathi Reddy Pedda Muntala, *Advanced API Security Techniques and Service Management*, 3 INT'L J. EMERGING RSCH. ENG'G & TECH. 63, 67–68 (2022).

42. Brown, *supra* note 40.

independently.[43] Customers are still responsible for their data and content, but Google encrypts all customer content at rest by default.[44] Thus, data stored with Google Cloud can be encrypted such that the government or third parties do not have access to its contents.[45] Like the other two compute providers, Google Cloud has a confidential computing offering that protects the confidentiality of cloud-based data by encrypting it while it is being processed, thereby securing AI workloads on its machines.[46] The hardware-based encryption technique ensures that Google Cloud cannot access the data as it is being processed. Google Cloud's documentation specifically highlights confidential AI as a use case for its confidential computing offering.[47]

2. Policy Safeguards Limit Provider Access to Customer Data.

All three compute providers in their customer data policies cabin the scenarios in which they may access customer data. Amazon does not access or use customer content for any purpose without the customer's consent, and it only discloses customer content when required to do so to comply with the law or a binding order of a government body.[48] Microsoft Azure similarly denies its personnel access to customer data by default.[49] It explicitly states that it does not build backdoors into any part of its product or provide any government with encryption keys or the ability to break its encryption.[50] It only responds to requests for information that go through a valid legal process, e.g., a valid subpoena or court order.[51] Google's terms are less explicit with respect to its

43. Nancy Liu, *Google Cloud CISO Contrasts Shared Fate vs. Shared Responsibility Models*, SDXCENTRAL (Mar. 22, 2024), https://www.sdxcentral.com/articles/interview/google-cloud-ciso-contrasts-shared-fate-vs-shared-responsibility-models/2024/03/ [https://perma.cc/H4UQ-JJP3]; *see also Shared Responsibilities and Shared Fate on Google Cloud*, GOOGLE CLOUD, https://cloud.google.com/architecture/framework/security/shared-responsibility-shared-fate [https://perma.cc/9GDG-G3E7].

44 . *Default Encryption at Rest*, GOOGLE CLOUD, https://cloud.google.com/docs/security/encryption/default-encryption [https://perma.cc/V26P-BSTM].

45. Google also explicitly states that it will not give any government entity "backdoor" access to the customer's data or to its servers storing that data in its HIPAA overview guide. *Google Cloud HIPAA Overview Guide*, GOOGLE CLOUD 1, 16–17 (May 2023), https://services.google.com/fh/files/misc/hipaa_overview_guide_googlecloud_whitepaper.pdf [https://perma.cc/WTY3-P6FL].

46 . *Confidential Computing*, GOOGLE CLOUD, https://cloud.google.com/security/products/confidential-computing [https://perma.cc/9LSJ-CHQ2].

47. *Confidential Computing for Data Analytics, AI, and Federated Learning*, GOOGLE CLOUD, https://cloud.google.com/architecture/confidential-computing-analytics-ai [https://perma.cc/HY9M-X3Q4].

48. *AWS Data Processing Addendum*, AMAZON WEB SERVS., https://d1.awsstatic.com/legal/aws-dpa/aws-dpa.pdf [https://perma.cc/6NQ5-BPND].

49. *Azure Customer Data Protection*, MICROSOFT (Sep. 29, 2024), https://learn.microsoft.com/en-us/azure/security/fundamentals/protection-customer-data [https://perma.cc/6TDM-ZK6Q].

50. MICROSOFT, *supra* note 33.

51. *Id.*

access to customer data. Per its terms, Google may process customer data in accordance with its agreement with the customer and applicable law to "provide, secure, and monitor" its services and technical support service.[52] However, the terms also require it to ensure compliance with its security measures, which include encryption of customer data.[53]

These technical and contractual guarantees underscore the extent to which the leading compute providers have designed their systems and built their customer commitments around strong data confidentiality. Through encryption, confidential computing, and narrowly circumscribed access policies, compute providers have deliberately limited their own visibility into customer data. Consequently, any compute monitoring provision that requires compute providers to access or disclose detailed information about AI training runs or data could conflict with both the technical design of these systems and the assurances the compute providers have given their customers. In fact, compliance with such a provision may be technically infeasible without compelling providers to develop new mechanisms, e.g., decryption software or backdoors, that undermine the confidentiality assurances they offer. Users of compute providers may argue that monitoring of their cloud-based data violates their expectation of privacy in their data under the Fourth Amendment, while compute providers may challenge compelled decryption as compelled speech under the First Amendment.

B. Fourth Amendment Challenges

Requiring providers to disclose information pertaining to users and their computational usage may lead to Fourth Amendment unreasonable search challenges based on users' reasonable expectation of privacy in their cloud-based data. After all, users may reasonably expect their cloud-based data, including data used to train AI models, to be kept private in light of the encryption technologies and privacy assurances described in the preceding section. Indeed, mandated monitoring and reporting by technology companies in other contexts, like required reporting of child sexual abuse material ("CSAM"), have raised such challenges.[54] This Section assesses whether requiring compute providers to report information about users and their AI training runs to the government constitutes a search under the Fourth Amendment, and if

52. *Cloud Data Processing Addendum (Customers)*, GOOGLE CLOUD (Aug. 21, 2025), https://cloud.google.com/terms/data-processing-addendum [https://perma.cc/L8RE-JGVF].

53. *Id.*

54. Federal law requires online platforms to report any discovered CSAM to the National Center for Missing and Exploited Children. 18 U.S.C. § 2258A. Courts have since addressed Fourth Amendment issues relating to CSAM reporting, such as whether the technology platforms are government agents and whether the government needs a warrant to review materials previously searched and flagged by the platform. *See, e.g.*, United States. v. Miller, 982 F.3d 412 (6th Cir. 2020); United States v. Wilson, 13 F.4th 961 (9th Cir. 2021).

so, whether such a search is constitutionally permissible under the private search, third-party, and consent doctrines.

1. Compute Providers May Be Deemed Government Agents If They Lack an Independent Purpose for Monitoring Customer Data.

To trigger Fourth Amendment protection, the conduct at issue must constitute state action.[55] As such, a private party could conduct a search that would be unconstitutional if conducted by the government.[56] Under the private search doctrine, a purely private search, even one that would be unconstitutional if performed by the government, does not implicate the Fourth Amendment if the government merely repeats or relies on what the private party already discovered.[57] But if the private party conducts a search while intending primarily to assist the government, it becomes a government agent, and the search must comply with the Fourth Amendment.[58] Thus, whether compute providers are subject to Fourth Amendment constraints hinges on whether they act independently or are co-opted into executing a government mandate.

Absent a clear test from the Supreme Court, lower courts have created their own tests for determining when private actors become government agents for Fourth Amendment purposes.[59] The frameworks developed by the lower courts thus far focus on the private actor's subjective intent, not solely on the degree of governmental control over the search. Thus, these frameworks are hard to apply in practice and introduce uncertainty for technology companies, who may be motivated by

55. United States v. Jacobsen, 466 U.S. 109, 113 (1984).

56. *Wilson*, 13 F.4th at 967.

57. *See id.*; United States v. Ackerman, 831 F.3d 1292, 1295 (10th Cir. 2016).

58. *See Ackerman*, 831 F.3d at 1300–01.

59. Jeff Kosseff, *Online Service Providers and the Fight Against Child Exploitation: The Fourth Amendment Agency Dilemma*, LAWFARE: THE DIGIT. SOC. CONT. 1, 3–4 (Jan. 2021), https://s3.documentcloud.org/documents/20458337/online-service-providers-and-child-exploitation.pdf [https://perma.cc/5YSZ-CX9S]. The Fifth, Ninth, and Eleventh Circuits apply the *Walther* framework, which considers two factors: (1) the government's knowledge of and acquiescence in the private party's conduct; and (2) the intent of the party performing the search to assist the government rather than to further its own ends. *See* United States v. Paige, 136 F.3d 1012, 1017–18 (5th Cir. 1998); United States v. Walther, 652 F.2d 788, 792 (9th Cir. 1981); United States v. Steiger, 318 F.3d 1039, 1045 (11th Cir. 2003). The First Circuit considers multiple factors: "the extent of the government's role in instigating or participating in the search, its intent and the degree of control it exercises over the search and the private party, and the extent to which the private party aims primarily to help the government or to serve its own interests." United States v. Pervaz, 118 F.3d 1, 6 (1st Cir. 1997). Finally, the Sixth and Eighth Circuits apply three tests, requiring satisfaction of at least one for the action at issue to constitute a government action: (1) whether the private party performed a public function (the "function" test); (2) whether the government compelled the private party's action (the "compulsion" test), and (3) whether the private party cooperated with the government (the "nexus" test). *Miller*, 982 F.3d at 422–23 (citing United States v. Ringland, 966 F.3d 731, 735 (8th Cir. 2020)).

both business interests (e.g., protecting infrastructure from misuse[60]) and legal obligations to assist law enforcement.[61] In *United States v. Bebris*,[62] Facebook's conclusory statement that it had "an independent business purpose" for monitoring its platform for CSAM was sufficient to show it was not a government agent.[63] Additionally, Microsoft's stipulation explaining that the presence of CSAM could substantially harm its image and reputation in the marketplace was also deemed sufficient.[64] Compute providers could invoke similar rationales in court. In particular, if a user was developing a model in a way that violated the provider's terms, the provider could argue it was independently interested in enforcing its terms or acceptable use policy.[65]

On the other hand, if compute providers are compelled to access encrypted content or peer into customer data in ways they do not currently do — and that conflicts with their privacy-focused design — they may be seen as acting solely at the government's behest.[66] For example, if providers monitor training data or other model data rather than compute usage alone, to comply with a statutory requirement, courts may find they are performing government-initiated searches subject to Fourth Amendment requirements.

The specific conduct that compute providers would be required to perform also matters. In the context of CSAM monitoring, several federal courts of appeals have held that Internet service providers are not government actors for Fourth Amendment purposes when they voluntarily search for CSAM on their platforms, despite their statutory obligation to report CSAM.[67] In *United States v. Miller*, for instance, the Sixth Circuit distinguished Google's statutory obligation to report *known* CSAM from a law that would compel or encourage Google to affirmatively monitor, search, screen, or scan its customers' files for

---

60. *See, e.g.*, Google Threat Intelligence Group, *Adversarial Misuse of Generative AI*, GOOGLE CLOUD: BLOG (Jan. 29, 2025), https://cloud.google.com/blog/topics/threat-intelligence/adversarial-misuse-generative-ai [https://perma.cc/B9M9-HSEU] (explaining how Google investigates activity associated with threat actors, including the misuse of generative AI or LLMs, to protect its products and users against cyber threats).

61. Kosseff, *supra* note 59, at 5.

62. 4 F.4th 551 (7th Cir. 2021).

63. *Id.* at 561.

64. *Id.*

65. *See* Orin S. Kerr, *Terms of Service and Fourth Amendment Rights*, 172 U. PA. L. REV. 287, 322 (2024).

66. A private party with some private or economic purpose may still qualify as a government agent. *See* United States v. Ackerman, 831 F.3d 1292, 1303 (10th Cir. 2016). Previous cases found that providers were not government agents when their actions were motivated by their "wholly private interests," United States v. Keith, 980 F. Supp. 2d 33, 40 (D. Mass. 2013), or "an independent business purpose." *Bebris*, 4 F.4th at 561. But there is little case law addressing the gray area where a decision to conduct a search is made by many employees who might have different motivations, some law enforcement-related and others business-related. Kosseff, *supra* note 59, at 5.

67. *See, e.g.*, United States v. Miller, 982 F.3d 412, 424 (6th Cir. 2020); *Bebris*, 4 F.4th at 562.

CSAM.[68] No statute compelled the specific conduct at issue — hash-value matching — so Google was not a government agent under the "compulsion" test.[69] Similarly, a compute monitoring provision like the draft rule issued by the U.S. Department of Commerce would not convert a compute provider into a government agent. The rule mandates reporting of large AI model training when the compute provider has knowledge of a foreign person transacting with the provider to train a large AI model for potentially malicious uses, but does not require proactive monitoring, searching, or scanning of model usage.[70] The more a compute monitoring provision requires, however, the harder it is for a compute provider to argue that its conduct was motivated entirely by its business interests and not to help law enforcement.

Even if compute providers are not deemed government agents for Fourth Amendment purposes, under the private search doctrine, the government must not exceed the scope of the earlier private search.[71] That is, subsequent review or searching by the government should not disclose more than what the private party already discovered.[72] If a compute provider implemented an automated mechanism for flagging potentially malicious training runs without an actual employee reviewing and confirming that the training run was indeed used to develop a model for harmful purposes,[73] warrantless review of the underlying data by the government could exceed the scope of the compute provider's original search because it could reveal new, protected information.[74] Thus, to avoid running afoul of the Fourth Amendment, a compute provider who wishes to comply with monitoring provisions must walk a fine line between conducting its own review of flagged

68. 982 F.3d at 424.

69. *Id.*

70. Taking Additional Steps to Address the National Emergency with Respect to Significant Malicious Cyber-Enabled Activities, 89 Fed. Reg. 5698, 5706 (proposed Jan. 29, 2024) (to be codified at 15 C.F.R. pt. 7).

71. United States v. Ringland, 966 F.3d 731, 736 (8th Cir. 2020).

72. *Id.* at 737.

73. A separate concern, outside the scope of this Note, is how a compute provider (or government) could reliably ascertain whether a model was being developed for potentially malicious purposes. *See* Bommasani et al., *supra* note 13.

74. In the CSAM scanning context, the Ninth and Tenth Circuits have concluded that the government exceeded the original private search where the private party automatically flagged files as CSAM but did not actually confirm whether the files were CSAM first. *See* United States v. Wilson, 13 F.4th 961, 976 (9th Cir. 2021) (where the government viewed the flagged materials but no one at Google had previously viewed them); United States v. Ackerman, 831 F.3d 1292, 1306–07 (10th Cir. 2016) (where the government agent viewed an e-mail that AOL had not previously examined). The Fifth and Sixth Circuits have held that subsequent review by the government did not exceed the scope of the private search, however. *See* United States v. Reddick, 900 F.3d 636, 639 (5th Cir. 2018) (characterizing the opening of the file by the government as merely confirmatory); *Miller*, 982 F.3d at 429–30 (relying on the lower court's finding that Google's flagging system was sufficiently reliable to conclude that the later government search revealed nothing more than what private parties had already seen).

training instances — to make it more likely the private search doctrine applies — while maintaining sufficiently independent business reasons for that review, other than to help the government.[75]

Currently, leading compute providers focus on monitoring resource usage for performance, resource management, and security purposes — but there is no public indication that they delve into the specific details or content of models being trained using their infrastructure.[76] If the compute providers were subjected to a compute monitoring rule that required them to actively monitor their platforms for training runs for certain malicious purposes, this lack of (publicly known) monitoring of model training details by compute providers would cut against a finding that they had an independent private purpose. The compute providers or the government could cite to the dominant compute providers' acceptable use policies — all of which contemplate and prohibit some malicious uses of their platforms — as evidence that compute monitoring would serve an independent business purpose, thereby barring applicability of the Fourth Amendment.[77] And as mentioned earlier, a rule that required only disclosure of training instances known by the provider to be for malicious purposes, as opposed to active and affirmative data collection, would likely be consistent with compute providers' existing terms and not be viewed as compulsion. But it would be difficult to argue that more detailed compute monitoring, which did not exist beforehand and was created in response to a new law, was born out of independent, purely private reasons. Absent an independent rationale, the compute provider may be found a government agent, and the Fourth Amendment would apply to its searches of customer data.

In short, compute providers will likely avoid being deemed government agents when their monitoring is limited to known conduct that

---

75. *See* Kosseff, *supra* note 59, at 13–14.

76. For example, under Amazon's Shared Responsibility Model, AWS is responsible for security at the software and hardware level, which includes monitoring compute, storage, databases, and networking. *See* AMAZON WEB SERVS., *supra* note 37. But application-level security and data fall under the customer's purview, *id.*, such that AWS does not access or use customer data without the user's consent by default "except as required . . . to comply with law." *Data Privacy Center*, AMAZON WEB SERVS., https://aws.amazon.com/compliance/data-privacy [https://perma.cc/EF3V-D8WU]. As mentioned in Section III.A., such monitoring could be technically infeasible for compute providers' privacy-first or encrypted offerings.

77. The AWS Acceptable Use Policy prohibits usage of AWS services "for any illegal or fraudulent activity; to violate the rights of others; to threaten, incite, promote, or actively encourage, violence, terrorism, or other serious harm," and reserves to AWS the right to investigate any suspected violation of its policy. *AWS Acceptable Use Policy*, AMAZON WEB SERVS., https://aws.amazon.com/aup/ [https://perma.cc/X2UZ-AQL2]. The Google Cloud Platform Acceptable Use Policy likewise prohibits customers from using the services to engage in illegal activity, violate others' legal rights, distribute viruses or other items of a destructive nature, or for any unlawful or invasive purpose. *Google Cloud Acceptable Use Policy*, GOOGLE CLOUD, https://cloud.google.com/terms/aup [https://perma.cc/9ZN6-NF7U].

violates their usage policies. However, broader monitoring obligations that require access to encrypted content could place them at risk of classification as government actors, triggering Fourth Amendment protections.

2. Technical Protections and the Ubiquity of Cloud-Based Data Can Support a Reasonable Expectation of Privacy.

Under the Fourth Amendment, a search occurs when the government violates a person's reasonable expectation of privacy.[78] The person must have an actual, subjective expectation of privacy, and that expectation must also be reasonable.[79] At a minimum, courts have found that applying security measures like password protection and encryption to data evinces a subjective expectation of privacy.[80] For customers of compute providers, who may rely on encryption and confidential computing services, that requirement is straightforwardly satisfied.

The harder question is whether those users have a reasonable expectation of privacy in cloud-stored data. Courts have not addressed, let alone reached a consensus, on that specific issue.[81] But recent case law and commentary concerning privacy in the digital age offer some guidance.

Many commentators who contend that users have a reasonable expectation of privacy in encrypted data analogize encryption to putting the content at issue into an opaque container or locking it with a key.[82] In the physical world, when assessing the contents of tangible containers, courts consider the nature of the container and any reasonable steps the owner took to conceal its contents.[83] Analogizing to the digital context, some courts have protected the "virtual container" in which digital data resides.[84] And so the argument goes: just as physical concealment efforts, like locking a briefcase, have been found to create a reasonable expectation of privacy, virtual concealment efforts, like encrypting or

78. Katz v. United States, 389 U.S. 347, 361 (1967) (Harlan, J., concurring).

79. *Id.*

80. *See, e.g.*, United States v. Yudong Zhu, 23 F. Supp. 3d 234, 238 (S.D.N.Y. 2014); United States v. Ziegler, 474 F.3d 1184, 1189 (9th Cir. 2007).

81. *See* Johnson v. VanderKooi, 983 N.W.2d 779, 793 (Mich. 2022) (Welch, J., concurring); *see also* Candice Gliksberg, *Decrypting the Fourth Amendment: Applying Fourth Amendment Principles to Evolving Privacy Expectations in Encryption Technologies*, 50 LOY. L. REV. 765, 790 (2017) (stating that "whether encryption itself triggers Fourth Amendment protection has not yet been addressed by the Supreme Court").

82. Orin S. Kerr, *The Fourth Amendment in Cyberspace: Can Encryption Create a "Reasonable Expectation of Privacy?"*, 33 CONN. L. REV. 503, 504 (2001).

83. David A. Couillard, Note, *Defogging the Cloud: Applying Fourth Amendment Principles to Evolving Privacy Expectations in Cloud Computing*, 93 MINN. L. REV. 2205, 2209 (2009).

84. *Id.* at 2219.

password-protecting a virtual container should likewise create a reasonable expectation of privacy in the concealed data.[85] Indeed, one lower court has likened encrypted files to items placed in opaque containers.[86] Another has deemed the use of passwords and encryption to support a finding of a reasonable expectation of privacy in the context of a government search of a password-protected, encrypted laptop.[87] Given that more data is now stored in the cloud than on local devices, the logic of these decisions applies with even greater force to cloud-based data secured by encryption.[88]

Not all commentators agree. For instance, Orin Kerr has argued that encryption cannot create a reasonable expectation of privacy, in part because encryption is similar to encoding or translating text, neither of which was historically protected by the Fourth Amendment.[89] In response, commentators have distinguished encryption, which uses a unique, indecipherable key that is shared between at most two users, from encoded communications or communications in a foreign language.[90] Such a key is generally not guessable or readily ascertained from general intelligence and education, in contrast to discovering an idiosyncratic code word, applying a simple substitution cipher, or translating from a known language.[91] Encryption, in their view, is not just obfuscation, but also a method of exclusion.

Moreover, recent cases from the Supreme Court, decided after the publication of Kerr's argument in 2001, have emphasized the importance of the nature of the technology and how it is used in society. In *Kyllo v. United States*,[92] the Court considered whether thermal imaging of a home constituted a search under the Fourth Amendment. Writing for the majority, Justice Scalia considered "what limits there are upon this power of technology to shrink the realm of guaranteed privacy" and observed that any rule the Court adopted "must take account of more sophisticated systems that are already in use or in

85. *See id.* at 2225.

86. United States v. Kim, 677 F. Supp. 2d 930, 943 (S.D. Tex. 2009).

87. United States v. Yudong Zhu, 23 F. Supp. 3d 234, 238–39 (S.D.N.Y. 2014). However, the court concluded that the government search at issue did not violate the Fourth Amendment on third party doctrine grounds. *Id.* at 240–41.

88. As explained below, the Supreme Court has considered the volume of potentially sensitive data in Fourth Amendment cell phone cases in determining whether a party had a reasonable expectation of privacy in the data at issue. *See infra* note 104; Carpenter v. United States, 585 U.S. 296, 320 (2018) (considering "the deeply revealing nature of CSLI, its depth, breadth, and comprehensive reach").

89. Kerr, *supra* note 82, at 506.

90. Gliksberg, *supra* note 81, at 778. Commentators have also pointed out that the searches in the cases Kerr relies on to make his argument were held reasonable on other grounds, e.g., because the items at issue were abandoned. *See, e.g.*, Stephen E. Henderson, *Nothing New Under the Sun? A Technologically Rational Doctrine of Fourth Amendment Search*, 56 MERCER L. REV. 507, 532 n.135 (2005).

91. *See* Gliksberg, *supra* note 81, at 778.

92. 533 U.S. 27 (2001).

development."[93] While the decision hinged primarily on the home being a protected zone with a minimum reasonable expectation of privacy,[94] it was an early indication that the Court was willing to consider technology's relevance in determining what constitutes a Fourth Amendment search.

Subsequent decisions confirmed this shift. In *United States v. Jones*,[95] a case involving location monitoring of a vehicle through a physical GPS tracking device, the concurring Justices considered the nature of long-term location monitoring as part of their inquiry into whether there was a reasonable expectation of privacy.[96] Justice Alito, writing for four Justices, opined that technology can alter expectations of privacy, especially as new technologies "provide increased convenience or security at the expense of privacy" and offer less of a meaningful choice to the public.[97] He highlighted the ubiquity of location monitoring — both of vehicles and of cell phone users — as something that would alter the average person's expectation of privacy in his or her location.[98] Justice Sotomayor, in her concurrence, agreed that basic Fourth Amendment premises may need to be reconsidered in the digital age, "in which people reveal a great deal of information about themselves to third parties in the course of carrying out mundane tasks."[99]

These concerns regarding how technology would alter expectations of privacy came to fruition in cases involving cell phones. In *Riley v. California*,[100] the Court held that the contents of a cell phone could not be searched without a warrant — even when the phone is seized incident to arrest[101] — because of qualitative and quantitative differences between cell phones and other physical containers.[102] The *Riley* Court emphasized that cell phones carry immense storage capacity and large volumes of data, and the data stored often contains sensitive personal information.[103] And in *Carpenter v. United States*,[104] the Court emphasized the pervasiveness and ubiquity of cell phones, observing that cell-site location information ("CSLI") collected from phones is not truly voluntarily exposed given the necessity of carrying phones to participate in modern society and the lack of affirmative user action needed for such tracking.[105] Thus, the Court has shown a willingness to

93. *Id.* at 34, 36.
94. *Id.* at 28.
95. 565 U.S. 400 (2012).
96. *Id.* at 415 (Sotomayor, J., concurring); *id.* at 430 (Alito, J., concurring).
97. *Id.* at 427 (Alito, J., concurring).
98. *Id.* at 428–29.
99. *Id.* at 417 (Sotomayor, J., concurring).
100. 573 U.S. 373 (2014).
101. *Id.* at 373–74.
102. *Id.* at 375.
103. *Id.*
104. 585 U.S. 296, 311 (2018).
105. *Id.* at 298.

consider technology's role in society when determining whether a reasonable expectation of privacy exists.

These precedents suggest that cloud computing ought to be analyzed in light of its scale and societal use. Compute providers offer massive data storage and processing capabilities, and a growing number of users entrust them with sensitive personal or proprietary information, with the expectation that their information will be kept private by technical and contractual guarantees.[106] And unlike CSLI, which is passively generated, use of confidential computing or encryption may reflect active or intentional concealment, strengthening the expectation of privacy.[107] Training data stored in encrypted environments reflects a deliberate choice regarding security and confidentiality.

Contrary to Kerr's early prediction that courts would "apply the principles already established by the Fourth Amendment in the physical world to the Internet,"[108] recent cases have shown a growing willingness to treat digital data as qualitatively and quantitatively distinct with respect to scale, ubiquity, and sensitivity. At the same time, commentators have relied on physical analogies to illustrate that encryption may be the only available method of securing digital content — akin to a lock or opaque container in the physical world.[109] As such, encryption of cloud-based data may be a reasonable step taken to ensure privacy that supports a reasonable expectation of privacy.

3. The Third-Party Doctrine May Limit Privacy Protections for Compute Usage Metrics but Not for Encrypted Content.

Even if the nature of technologies like encryption could support a reasonable expectation of privacy in cloud-based data, the third-party doctrine still poses an obstacle. That doctrine provides that information voluntarily revealed to third parties is not protected by the Fourth Amendment.[110] As applied to data stored by compute providers, the third-party doctrine is not likely to extinguish a reasonable expectation of privacy in all cloud-based data. However, the Fourth Amendment likely does not protect non-content data, such as compute metrics.

As described earlier, the Supreme Court indicated in *Riley* and *Carpenter* its willingness to distinguish disclosure of digital data from other

106. *See supra* Section III.A

107. United States v. Soybel, 13 F.4th 584, 593 (7th Cir. 2021) (explaining that while CSLI is collected passively whenever a cell phone is powered on, a user must act affirmatively to generate IP data). By deciding to use certain services or enable certain features, a cloud compute user likewise acts affirmatively, thereby evincing an expectation of privacy.

108. Kerr, *supra* note 82, at 532.

109. Couillard, *supra* note 83, at 2234.

110. Smith v. Maryland, 442 U.S. 735, 743–44 (1979) ("[A] person has no legitimate expectation of privacy in information he voluntarily turns over to third parties.") (collecting cases).

third-party doctrine cases. In *Riley*, the Court emphasized the intrusiveness of police access to cloud-based data via a cell phone by analogizing the cell phone to a key and the cloud-based data to a house.[111] The cell phone was like a key in a suspect's pocket that would allow law enforcement to unlock and search the suspect's house (i.e., the suspect's data).[112] In *Carpenter*, the Court highlighted the ubiquity and immense storage capacity of cell phones.[113]

Cloud computing parallels and exceeds the scenarios in those two cases. Users routinely entrust cloud computing platforms with vast quantities of data, including proprietary training data and private models.[114] While at first glance, training data may not appear as inherently personal as the browser history on one's cell phone, such datasets often contain identifiable or confidential information, particularly in domains like health and finance.[115] As cloud computing becomes more omnipresent, the disclosure of data stored on compute platforms becomes less voluntary under the third-party doctrine.[116]

Nevertheless, the third-party doctrine continues to apply to certain types of non-content data, i.e., transactional records. In *Smith v. Maryland,*[117] the Court distinguished the actual contents of telephonic communications from numbers that were dialed and conveyed to a telephone company, holding that the latter was not protected because the company recorded that information for "legitimate business purposes."[118] In the digital context, multiple circuits have explicitly held that there is no reasonable expectation of privacy in an IP address because the IP address itself conveys no substantive information.[119] Similarly, the to/from addresses on e-mails have also been considered transactional data,[120] in contrast to the contents of the e-mail communications. By analogy, transactional data exchanged with compute providers, such as compute purchased, usage logs, or resource provisioning details, will likely fall outside Fourth Amendment

111. Riley v. California, 573 U.S. 373, 397 (2014).

112. *Id.*

113. Carpenter v. United States, 585 U.S. 296, 305 (2018) (quoting *Riley,* 573 U.S. at 393).

114. Chris Arkenberg et al., *Taking Control: Generative AI Trains on Private, Enterprise Data*, DELOITTE (Nov. 29, 2023), https://www.deloitte.com/us/en/insights/industry/technology/technology-media-and-telecom-predictions/2024/tmt-predictions-enterprise-ai-adoption-on-the-rise.html [https://perma.cc/7B4Z-QGGC].

115. *Id.*; Ghiath Shabsigh & El Bachir Boukherouaa, *Generative Artificial Intelligence in Finance: Risk Considerations*, IMF 5–6 (Aug. 2023), https://www.imf.org/-/media/files/publications/ftn063/2023/english/ftnea2023006.pdf [https://perma.cc/QM7B-AWKN].

116. U.S. Dep't of Just. v. Ricco Jonas, 24 F.4th 718, 739 (1st Cir. 2022).

117. 442 U.S. 735 (1979).

118. *Id.* at 743.

119. *See* United States v. Jean, 207 F. Supp. 3d 920, 932 (W.D. Ark. 2016) (collecting cases), *aff'd*, 891 F.3d 712 (8th Cir. 2018).

120. Couillard, *supra* note 83, at 2228.

protection, as these metrics are used for routine business purposes.[121] Encrypted, cloud-based content data, however, would remain protected.[122]

4. Whether Agreement to Terms of Service Constitutes Voluntary Consent Depends on the Specificity and Clarity of Those Terms.

Regardless of whether users have a reasonable expectation of privacy in their cloud-based data, compute providers may try to obtain consent for searches through their terms of service. If the subject of a search provides voluntary consent, then neither a warrant nor probable cause is required for the search to be permissible.[123] However, as illustrated below, courts are divided on whether agreeing to a compute provider's terms of service amounts to valid consent, and more broadly, on whether such agreements can affect a user's expectation of privacy. The answer is highly fact-specific and turns on both the language of the terms and the user's understanding of them.[124]

Through compiling and classifying recent cases, Kerr has identified a divide in the case law: some courts hold that terms of service can reduce or eliminate Fourth Amendment protections, while others hold that users have a reasonable expectation of privacy beyond the terms of service.[125] Many of the cases treating terms of service as controlling have involved cloud providers scanning the contents of user account files for known images of CSAM.[126] For example, in *United States v. Bohannon*,[127] the court held that Microsoft did not violate the Fourth Amendment by scanning folders on its OneDrive cloud service.[128] The court found that the defendant consented to the search by agreeing to Microsoft's terms of service,[129] which allowed Microsoft to access and disclose his personal data and private folders stored on the cloud, "so long as Microsoft believed in good faith that doing so was necessary to enforcing its terms of service."[130] The defendant's agreement to the terms of service thus precluded any reasonable expectation of privacy

121. For example, Microsoft describes its non-content data as including basic subscriber information, such as address and payment information, and its content data as what its customers create and store through its services, including content on its cloud offerings. MICROSOFT, *supra* note 34.

122. Couillard, *supra* note 83, at 2237.

123. Schneckloth v. Bustamonte, 412 U.S. 218, 219 (1973) (citing Davis v. United States, 328 U.S. 582, 593–94 (1946); Zap v. United States, 328 U.S. 624, 630 (1946)).

124. Kosseff, *supra* note 59, at 15 & n.65.

125. Kerr, *supra* note 65, at 294.

126. *Id.* at 296.

127. 506 F. Supp. 3d 907 (N.D. Cal. 2020).

128. *Id.* at 915.

129. *Id.* at 913.

130. *Id.* at 915.

in his cloud-based data.[131] Along similar lines, the Minnesota Court of Appeals ruled in *State v. Pauli* that the defendant had no Fourth Amendment rights in files he remotely stored on his Dropbox account because Dropbox's terms of service undermined his reasonable expectation of privacy.[132] The Court reasoned that because the defendant breached the terms of service by uploading CSAM to Dropbox, Dropbox could then review his content for compliance, eliminating any reasonable expectation of privacy.[133]

However, other courts have declined to treat terms of service as dispositive. Instead, these courts have relied on physical analogies, observing that individuals "often have Fourth Amendment rights in physical spaces despite having granted rights of access to third parties."[134] In *United States v. Warshak*,[135] the Sixth Circuit held that an internet provider's terms of service did not diminish defendants' expectations of privacy in their email contents, even though the terms stated that the provider could access and use individual user information.[136] The Sixth Circuit reasoned that the provider's retention of the right of access did not diminish the reasonableness of the defendants' trust in the privacy of his emails.[137] Though it was possible a sweeping user agreement could defeat a reasonable expectation of privacy, the Sixth Circuit was skeptical that would be the case in most situations.[138] Likewise, the district court in *United States v. Irving* held that the defendant had a reasonable expectation of privacy in his Facebook account contents because Facebook's terms did not explicitly eliminate user rights and because Facebook had not known that there was a breach of its terms to begin with.[139] Notably, Facebook's terms of service stated that the user owned all content and information and could control sharing it and did not explicitly note that Facebook would monitor the user's account for illegal activities and report them to law enforcement.[140] The district court in *United States v. DiTomasso* took a mixed approach.[141] Though it held that agreeing to terms of service would not eliminate a reasonable expectation of privacy, it also held that using a service would constitute consent if the terms sufficiently put the user on notice that the

131. *Id.* at 916 n.5.
132. State v. Pauli, No. A19-1886, 2020 WL 7019328, at *8–9 (Minn. Ct. App. Feb. 24, 2021), *aff'd on other grounds*, 979 N.W.2d 39 (Minn. 2022).
133. *Id.* at *7.
134. Kerr, *supra* note 65, at 300.
135. 631 F.3d 266 (6th Cir. 2010).
136. *Id.* at 286.
137. *Id.* at 287.
138. *Id.* at 286.
139. 347 F. Supp. 3d 615, 623 (D. Kan. 2018).
140. *Id.*
141. 56 F. Supp. 3d 584 (S.D.N.Y. 2014), *aff'd on other grounds*, 932 F.3d 58 (2d Cir. 2019).

provider would cooperate with law enforcement.[142] The court reached its decision by differentiating between one policy that included only a passing reference to law enforcement and another policy that made clear the provider would actively assist law enforcement.[143]

Finally, even where there is notice of potential government cooperation, the defendant's own familiarity and understanding of the policy presents an issue of fact. Consent must be voluntary,[144] yet it is well-documented that many users usually scroll past the terms of service.[145]

As applied to compute providers, this question remains open.[146] The terms of service for AWS, Microsoft Azure, and Google Cloud all specify that customers retain responsibility and control over their data.[147] No policy explicitly states that the provider will affirmatively monitor user content for law enforcement purposes or disclaims user privacy rights. Under the reasoning found in *Warshak*, *Irving*, and *DiTomasso*, the language in the terms of service may be insufficient to override a user's reasonable expectation of privacy in their cloud-based data.

C. First Amendment Challenges

Beyond the unsettled Fourth Amendment jurisprudence surrounding cloud-based data, compute monitoring also raises distinct First Amendment concerns. If compute providers are not only permitted to monitor certain user data but are also required to do so, they may be forced to develop decryption tools or backdoors. This raises the question of whether such a requirement constitutes impermissible compelled speech. Courts have generally upheld disclosure mandates where the information is purely factual[148] and the government interest is substantial.[149] But where compelled decryption is a *prerequisite* to that mandated disclosure, the constitutional question changes to whether the compute provider is being compelled to write code enabling that decryption.[150]

---

142. *Id.* at 597.

143. *Id.*

144. *See* Kosseff, *supra* note 59, at 15 n.65; Kerr, *supra* note 65, at 325 (arguing that courts should determine how terms of service "operate in real life" in the Fourth Amendment context).

145. Kerr, *supra* note 65, at 325.

146. Kerr, for his part, argues that the terms of service should be irrelevant to the issue of consent. *Id.* at 324.

147. *See supra* Section III.A.

148. *See* Am. Beverage Ass'n v. City & Cnty. of San Francisco, 916 F.3d 749, 756 (9th Cir. 2019).

149. *Id.* at 755 (holding that compelled truthful disclosure in commercial speech is constitutional so long as there is a substantial government interest).

150. *See infra* Section C.2

1. Disclosure of Routine Compute Data is Likely Permissible, But Subjective or Burdensome Reporting May Raise First Amendment Concerns.

The First Amendment's Free Speech Clause applies not only to laws prohibiting speech but also to laws compelling speech.[151] As a result, information disclosure requirements have been challenged as compelled speech.[152] Compute providers may similarly argue that mandated reporting to the government under a compute monitoring provision compels them to disclose information they would not otherwise share, potentially undercutting their business models, customer relationships, or public commitments to privacy.

Courts evaluate compelled disclosure requirements under varying levels of scrutiny based on the type of speech being compelled. Strict scrutiny applies by default,[153] but courts often apply lower levels of scrutiny to "purely factual and uncontroversial" commercial speech.[154] Under the Supreme Court's test in *Zauderer v. Office of Disciplinary Counsel*, the government may compel the disclosure of commercial speech so long as the disclosed information is "(1) purely factual, (2) noncontroversial, and (3) not unjustified or unduly burdensome."[155]

Whether reporting requirements under a compute monitoring rule meet the *Zauderer* standard depends heavily on the nature of the information to be reported. At one end of the spectrum, disclosure of transactional data, such as names, IP addresses, or payment details, is likely to be considered purely factual[156] and noncontroversial.[157] Providers

151. Wooley v. Maynard, 430 U.S. 705, 714 (1977) (explaining that the First Amendment protects "both the right to speak freely and the right to refrain from speaking at all") (citing Bd. of Educ. v. Barnette, 319 U.S. 624, 633–34 (1943)).

152. *See, e.g.*, Nat'l Inst. of Fam. & Life Advocs. v. Becerra, 585 U.S. 755, 778 (2018) [hereinafter *NIFLA*] (holding that a disclosure requirement imposed on pregnancy-related clinics was unduly burdensome and would chill protected speech); Memorandum of Points & Authorities in Support of Motion for Preliminary Injunction at 50, X Corp. v. Bonta, No. 23-cv-01939, 2023 WL 8948286, at *50 (E.D. Cal. Dec. 28, 2023) (arguing that a law requiring social media platforms to disclose information about their content moderation practices would chill exercise of editorial discretion).

153. R J Reynolds Tobacco Co. v. FDA, 96 F.4th 863, 876 (5th Cir. 2024).

154. Zauderer v. Off. of Disciplinary Couns., 471 U.S. 626, 651 (1985).

155. Am. Beverage Ass'n v. City & Cnty. of San Francisco, 916 F.3d 749, 756 (9th Cir. 2019) (citing *NIFLA*, *supra* note 152, at 768).

156. The Supreme Court has not expressly defined what constitutes "purely factual" information, but circuit courts have interpreted this element to turn on whether the information is factually accurate or undisputed. *See R J Reynolds*, 96 F.4th at 879 (inquiring whether the disclosures composed of only "information supported by facts" and "conclusions driven by those facts"); CTIA — Wireless Ass'n v. City of Berkeley, 928 F.3d 832, 846–47 (9th Cir. 2019) [hereinafter *CTIA II*] (inspecting sentence-by-sentence whether the text of the compelled disclosure is "literally true").

157. *See R J Reynolds*, 96 F.4th at 881 (interpreting *NIFLA* to hold that a factual statement is "controversial" under *Zauderer* where its truth is "not settled or is overwhelmingly

already collect this information in the ordinary course of business, and it lacks controversial content or disputed meaning. Similarly, logging that a training run occurred at a particular time using a certain quantity of compute resources is likely to be considered purely factual and uncontroversial.

However, reporting requirements that obligate compute providers to assess whether a model is potentially capable of being used in malicious activities introduce a degree of subjectivity. In making that assessment, providers may need to adjudge the purpose of the model being trained, which activities are malicious (absent a clear definition from the government), and whether the model is capable of such activities. Because providers may not be able to make these assessments with certainty,[158] these judgments may no longer be "purely factual" and therefore qualify as compelled speech that is subject to a more rigorous standard of scrutiny.[159]

Moreover, the burden imposed by such a provision may be nontrivial. If the government merely asks compute providers to report training runs above a certain compute threshold or information already collected for operational reasons, the requirement would likely be upheld as minimally burdensome.[160] But if compliance requires developing new mechanisms to inspect encrypted data, scrutinize user training activity, or infer model capabilities, the rule could be considered too burdensome, especially if the compute provider's commitments and services are built around minimal knowledge of or access to customer data.

Assuming that the reporting requirement involves only routine data disclosures, the requirement is unlikely to violate the First Amendment using *Zauderer* scrutiny, as a court would likely find it to be "'reasonably related' to a substantial government interest."[161] Multiple circuit

---

disproven or where the inherent nature of the subject raises a live, contentious political dispute"); *NIFLA*, *supra* note 152, at 756 (holding that information about state-sponsored abortion services was controversial because abortion was a controversial topic); Nat'l Ass'n of Wheat Growers v. Bonta, 85 F.4th 1263, 1278 (9th Cir. 2023) ("[S]aying that something is carcinogenic or has serious deleterious health effects — without a strong scientific consensus that it does — remains controversial.").

158. As explained earlier, customers of the major compute providers may opt to encrypt their data and models, making it difficult for providers to assess with any certainty the purpose of the model being trained, as those providers cannot access the underlying data. *See supra* Section III.A. Additionally, scholars have already shown how the amount of compute used to train a model does not map onto the model's purpose or capability for malicious use. *See infra* notes 170, 195.

159. *R J Reynolds*, 96 F. 4th at 876. See *supra* note 156 for the lower courts' definition of what is "purely factual."

160. For instance, tracking usage of compute, memory, and other resources is critical to AWS functionalities that allow it or users to manage servers and jobs. *See Compute*, AMAZON WEB SERVS., https://docs.aws.amazon.com/whitepapers/latest/aws-overview/compute-services.html [https://perma.cc/UL92-D3ZG].

161. Am. Beverage Ass'n v. City & Cnty. of San Francisco, 916 F.3d 749, 755 (9th Cir. 2019) (quoting Zauderer v. Off. of Disciplinary Couns., 471 U.S. 626, 651 (1985)).

courts have construed *Zauderer* broadly and have treated any legitimate state interest as sufficient.[162] Given that courts have found health, safety, and welfare to constitute substantial government interests,[163] they would likely find the government's interest in national security to be sufficient as well.[164]

Next, the government must show that the national security interest justifies the reporting requirements.[165] The harm sought to be addressed must be "potentially real[,] not purely hypothetical."[166] Ample evidence exists as to the risk of foreign entities using domestic compute products to engage in malicious cyber-enabled activities,[167] and previous administrations have put forth findings of increasing prevalence and severity of such activities.[168]

Finally, the reporting requirements must reasonably relate to the state interest.[169] Compute providers could contest the effectiveness of the reporting requirements in preventing foreign actors from engaging in malicious, cyber-enabled activities.[170] But under such a lenient standard of review, a court would likely defer to the government to

---

162. *See R J Reynolds*, 96 F.4th at 882–83 (joining the First, Second, Sixth, Ninth, and D.C. Circuits to apply *Zauderer* beyond consumer deception context).

163. *See, e.g.*, *CTIA II*, *supra* note 156, at 845 ("There is no question that protecting the health and safety of consumers is a substantial government interest.") (citing Posadas de Puerto Rico Assocs. v. Tourism Co. of Puerto Rico, 478 U.S. 328, 341 (1986)).

164. Exec. Order No. 14,110, 3 C.F.R. 657, 665 (2024) (grounding the reporting requirements in "the national emergency related to significant malicious cyber-enabled activities"). Preventing foreign governments and actors from undermining democratic institutions has been deemed a compelling interest under a strict scrutiny analysis. Washington Post v. McManus, 355 F. Supp. 3d 272, 298-99 (D. Md. 2019), *aff'd*, 944 F.3d 506 (4th Cir. 2019). As *Zauderer* articulates a lower standard, such a compelling interest would also be a substantial one. *See id.* at 297.

165. *R J Reynolds*, 96 F.4th at 884.

166. *NIFLA*, *supra* note 152, at 757 (quoting Ibanez v. Fla. Dep't of Bus. & Pro. Regul., Bd. of Acct., 512 U.S 136, 146 (1994)).

167. *See, e.g.*, *Disrupting Malicious Uses of AI by State-Affiliated Threat Actors*, OPENAI (Feb. 14, 2024), https://openai.com/blog/disrupting-malicious-uses-of-ai-by-state-affiliated-threat-actors [https://perma.cc/WYY7-KPA7].

168. *See* Exec. Order No. 13,694, 3 C.F.R. 297 (2015); Exec. Order No. 13,757, 3 C.F.R. 659 (2017); Exec. Order No. 13,984, 86 Fed. Reg. 6837 (Jan. 19, 2021).

169. *See R J Reynolds*, 96 F.4th at 885.

170. As some have already pointed out, malicious actors could evade detection through "structuring" or deconstructing compute-intensive projects into smaller sub-projects that would not trigger scrutiny from IaaS providers. Janet Egan & Lennart Heim, *Oversight for Frontier AI Through a Know-Your-Customer Scheme for Compute Providers*, CTR. FOR GOVERNANCE OF AI, at 13 (Oct. 25, 2023), https://cdn.governance.ai/Oversight_for_Frontier_AI_through_a_KYC_Scheme_for_Compute_Providers.pdf [https://perma.cc/WB6S-KLJD]. Alternatively, malicious foreign actors may further obfuscate their identities through lying, using stolen credentials or identity-proofing documents, hiding behind shell companies, or buying access to verified IaaS accounts. *See id.*; Info. Tech. & Innov. Found., Comment Letter on Exec. Order 13,984 (October 24, 2021), https://www.regulations.gov/comment/DOC-2021-0007-0005 [https://perma.cc/MP58-JVHX]; Info. Tech. Ind. Council, Comment Letter on Exec. Order 13,984 (October 25, 2021), https://www.regulations.gov/comment/DOC-2021-0007-0013 [https://perma.cc/YA7W-QY94].

implement a know-your-customer ("KYC") scheme[171] intended to scrutinize foreign customers' usage of domestic compute services to train large, dual-use models.[172]

In short, routine disclosures of factual, operational data are unlikely to violate the First Amendment under *Zauderer*. But where reporting requires subjective assessments or burdensome technical changes, courts may apply heightened scrutiny.

2. Requiring Providers to Write Decryption Tools or Mechanisms Likely Constitutes Compelled Speech.

While disclosing information from the cloud may not constitute compelled speech, requiring a company to create a backdoor to provide access to that information may be. This issue came to the forefront in 2016, when Apple fought a high-profile battle with the Federal Bureau of Investigation ("Bureau") after the Bureau demanded that Apple write software to bypass the security features of an iPhone used by a shooter.[173] A federal judge ordered Apple to help the Bureau, and Apple filed a motion to vacate that order.[174] The Bureau ended up accessing the phone through other methods, so the court never fully resolved the question, but the parties and amici briefs sparred on whether forcing Apple to build an encryption backdoor would constitute compelled speech in violation of the First Amendment.[175]

Given the proliferation of privacy-first offerings in the cloud, a compute monitoring provision that required compute providers to peer

171. "Know-Your-Customer" or KYC is a principle that requires businesses to verify the identity of their users as a prerequisite to granting them access to particular services. Egan & Heim, *supra* note 170, at 7. As implemented in the financial sector, a KYC scheme may also include risk assessments or profiles of users, increased due diligence for higher-risk users, ongoing monitoring of transactions to detect suspicious activity, and reporting of such activities to the government. *Id.*

172 . Executive Order 13,984 tasked the Secretary of Commerce with proposing regulations requiring IaaS providers to implement a KYC scheme to deter foreign malicious cyber actors' use of domestic IaaS products. Exec. Order No. 13,984 at 6837–38. The Biden Administration views these procedures as a way to limit foreign hackers' ability to abuse domestic cloud services. John Sakellariadis, *White House Moves to Push Foreign Hackers Out of U.S. Cloud*, POLITICO (Feb. 27, 2023, at 10:00 ET), https://www.politico.com/newsletters/weekly-cybersecurity/2023/02/27/white-house-moves-to-push-foreign-hackers-out-of-u-s-cloud-00084505 [https://perma.cc/D2KK-WZNB]. KYC schemes have been supported more broadly as a way to enable greater oversight of frontier AI model development and ensure domestic cloud services are not being abused. *See, e.g.*, Egan & Heim, *supra* note 170; Brad Smith, *How Do We Best Govern AI?*, MICROSOFT (May 25, 2023), https://blogs.microsoft.com/on-the-issues/2023/05/25/how-do-we-best-govern-ai/ [https://perma.cc/FP6G-QTAA].

173. Cynthia Brumfield, *US Department of Justice Push for Encryption Backdoors Might Run Afoul of First Amendment*, CSO (Nov. 4, 2019), https://www.csoonline.com/article/568029/us-department-of-justice-push-for-encryption-backdoors-might-run-afoul-of-first-amendment.html [https://perma.cc/Z9L2-QBUN].

174. *Id.*

175. *Id.*

into the data stored on their platforms raises similar issues around compelled speech. Many popular cloud services use encryption and allow for users to retain control over their own encryption keys, essentially barring the providers from accessing the stored data.[176] Just as accessing the underlying data is a prerequisite to determining if a model is being trained for malicious purposes, implementing a backdoor is a prerequisite to accessing the otherwise encrypted data.[177]

Since the late 1990s, courts have recognized that code can constitute protected speech under the First Amendment. The Ninth Circuit famously laid down the "code is speech" principle in *Bernstein v. United States Department of Justice*,[178] reasoning that computer programming languages, like mathematical equations or economic graphs, can be used to express ideas.[179] Though the Ninth Circuit eventually withdrew the decision, the idea persisted, and courts in subsequent cases have agreed that code constitutes a form of speech.[180] But not all code is protected speech. Many courts have recognized "a distinction between expressive computer code, which is constitutionally protected, and functional computer code."[181]

Courts have held that encryption software is expressive for First Amendment purposes. *Bernstein* did so in the academic context, holding that the software is expressive "in its source code form and as employed by those in the field of cryptography."[182] The Sixth Circuit in *Junger v. Daley*,[183] also addressing restrictions on encryption source code in the academic context, similarly held that the encryption software at issue was covered by the First Amendment because the code was "an expressive means for the exchange of information and ideas about computer programming."[184]

The line between expressive and purely functional code remains contested. Several commentators have argued that where encryption software is purely functional, compelled backdoors do not implicate

176. *See supra* Section III.A.

177. Because the government does not have access to the user's unique encryption key, it must instead seek exceptional access from the platform through a backdoor. Paula Bernardi & Celia Richardson, *What Is an Encryption Backdoor?*, INTERNET SOC'Y (May 2, 2025), https://www.internetsociety.org/blog/2025/05/what-is-an-encryption-backdoor/ [https://perma.cc/F4WD-ZLXV].

178. 176 F.3d 1132 (9th Cir. 1999), *reh'g granted*, *opinion withdrawn*, 192 F.3d 1308 (9th Cir. 1999).

179. *Id.* at 1141.

180. *See, e.g.*, Def. Distributed v. Platkin, 697 F. Supp. 3d 241, 257 (D.N.J. 2023) (collecting cases).

181. *Id.* at 257–58.

182. *Bernstein*, 176 F.3d at 1141 (footnote omitted).

183. 209 F.3d 481 (6th Cir. 2000).

184. *Id.* at 485.

First Amendment concerns.[185] But this argument is less persuasive when applied to compute providers. Encryption is not merely a technical feature; it is bound up with the compute provider's public-facing commitments to user privacy. The software and architecture reflect policy choices and views on privacy that are echoed in the providers' public statements.[186] Code that implements encryption or confidential computing is thus both functional *and* expressive. That dual character does not strip it of constitutional protection.[187]

Compelling a compute provider to undermine its assurances — effectively communicating a message it does not wish to convey — violates the core of the compelled speech doctrine.[188] The provider is forced to adopt the government's position on privacy and security rather than its own, forcing it into a position of hypocrisy that the doctrine is meant to prevent.[189] Were Amazon, Microsoft, and Google to have to create a backdoor for the government, their assurances to their customers would ring hollow.

A regulation compelling that message is thus content-specific, as it is based on "[agreement or] disagreement with the message it conveys,"[190] and may only be upheld if it is narrowly tailored to further a compelling state interest.[191] In contrast to a narrow order concerning the decrypting of a single phone at the center of a criminal investigation,[192] a broad encryption backdoor affecting all user accounts would likely not survive a strict scrutiny analysis.[193] Safety and security are quintessential compelling state interests, yet a mandate that compute providers add a backdoor to their encrypted offerings is far too broad and potentially ineffective. Access to a wide swath of cloud-based data is not narrowly tailored to achieve the government's interest, when only a fraction of the data may indicate a potential use or development of a large AI model for malicious purposes.[194] Even then, users with

185. *See, e.g.*, Mark C. Bennett, *Was I Speaking to You?: Purely Functional Source Code as Noncovered Speech*, 92 N.Y.U. L. REV. 1494, 1527 (2017); Kyle Langvardt, *Crypto's First Amendment Hustle*, 26 YALE J.L. & TECH. 130, 145, 149–50 (2023).

186. *See supra* Section III.A.

187. *Junger*, 209 F.3d at 484–85; *see also Bernstein*, 176 F.3d at 1141.

188. Allen Cook Barr, *Guardians of Your Galaxy S7: Encryption Backdoors and the First Amendment*, 101 MINN. L. REV. 301, 327 (2016) ("[E]ncryption backdoors would require the expression of a particular idea. . . .").

189. *See, e.g.*, Wooley v. Maynard, 430 U.S. 705, 717 (1977) (recognizing that being forced to platform or distribute a message is itself a First Amendment harm).

190. Turner Broad. Sys Inc. v. Fed. Commc'ns Comm'n, 512 U.S. 622, 642 (1994) (quoting Ward v. Rock Against Racism, 491 U.S. 781, 791 (1989)).

191. *See id.* at 662 (quoting *Ward*, 491 U.S. at 799).

192. This is how the government characterized the demand it placed on Apple, arguing that such a narrow order does not threaten the underlying concerns behind the compelled speech doctrine, like the vitality of public discourse. Bennett, *supra* note 185, at 1529–30.

193. *See* Barr, *supra* note 188, at 329.

194. The Draft Rule, for instance, encompasses general-purpose models trained on immense datasets. *See supra* note 16.

malicious intents have a variety of workarounds available to them, ranging from federated learning to performing their own encryption.[195] As in the Apple case, overbreadth and technical ineffectiveness weigh heavily against the constitutionality of a compelled decryption mandate.

IV. RECOMMENDATIONS

Compute governance schemes may raise constitutional concerns and pose some risk of abuse; however, this risk can be minimized through both policy and technological means. On the policy side, policymakers should consider substantive guardrails, such as limiting the type of information that the government can request and the entities subject to such regulations.[196] As explained earlier, limiting the information at issue to objectively defined compute usage metrics or transactional data can reduce friction with constitutional protections and the privacy technologies employed by compute providers.[197] Such information is likely to be readily available to providers without requiring them to surveil users' models and data. Because that information is not currently kept private from providers, users are less likely to be able to establish a reasonable expectation of privacy for that information. Additionally, courts would likely uphold reporting requirements of such information under the First Amendment, as providers would be expected to disclose purely factual and uncontroversial information, rather than be compelled to express a belief or opinion on what a user's model is capable of and being used for or to write software to access confidential user data.[198] However, while such policy guardrails may decrease the risk of overly intrusive surveillance of user data, the less information regulators receive, the less useful compute governance as a scheme may be. Ultimately, policymakers face a tradeoff between

195. Malicious actors may seek to avoid detection by "structuring" or federating compute-intensive projects into smaller, discrete sub-projects that fall below the reporting threshold. *See* Egan & Heim, *supra* note 170, at 13. Additionally, some attacks, like supply chain poisoning do not occur in the training process and instead rely on injecting a surgical change that is not detectable through compute. *See* Chris McGowan, *Generative AI and the Potential for Nefarious Use*, ISACA (Aug. 1, 2023), https://www.isaca.org/resources/news-and-trends/industry-news/2023/generative-ai-and-the-potential-for-nefarious-use [https://perma.cc/A8H7-APPG]. Finally, malicious models that do not require significant compute already exist. *See, e.g.*, Shenggan Cheng, Xuanlei Zhao, Guangyang Lu, Jiarui Fang, Zhongming Yu, Tian Zheng et al., *FastFold: Reducing AlphaFold Training Time from 11 Days to 67 Hours* (Feb. 5, 2023) (unpublished manuscript) (on file with arXiv), https://arxiv.org/abs/2203.00854 [https://perma.cc/DWX8-965Z].

196. Lennart Heim, Markus Anderljung & Haydn Belfield, *To Govern AI, We Must Govern Compute*, LAWFARE (Mar. 28, 2024, at 12:00 ET), https://www.lawfaremedia.org/article/to-govern-ai-we-must-govern-compute [https://perma.cc/PT6Q-CYXM].

197. *See supra* Section III.B.3.

198. *See supra* Section III.C.2

what is easy to measure (and does not raise privacy concerns) and what is most correlated with substantive risk.[199]

From a technological standpoint, privacy-preserving technologies may allow for oversight mechanisms that do not require access to the underlying cloud-based data. New privacy-enhancing technologies could enable regulators to receive only a limited set of information about whether the user is in compliance without revealing any other data.[200] On-chip firmware could also be limited to collect only information that is needed to verify compliance, thereby removing the need for compute providers to peer into the underlying data.[201] These on-chip mechanisms could allow users to measure and report the amount of compute they used as well as where the work was carried out, giving regulators a finer-grained picture without compromising the user's privacy.[202] If these technologies are adopted, courts may find that the burden of a compute governance scheme is less significant on providers. Moreover, because these technologies are privacy-preserving, regulators are less likely to repeat the encryption backdoor debate, as providers will no longer be compelled to surveil their users' data.

V. CONCLUSION

As policymakers confront the risks posed by large AI models, compute governance has emerged as an attractive regulatory lever. Yet, as this Note shows, overly intrusive compute monitoring provisions risk running afoul of the Fourth and First Amendments.

These legal challenges are not insurmountable but do require careful tailoring through both policy and technological means. By grounding compute governance in privacy-aware design and constitutional principles, policymakers can advance their security and safety objectives without sacrificing the civil liberties of their constituents.

---

199. Bommasani et al., *supra* note 13.

200 . *See* Girish Sastry, Lennart Heim, Haydn Belfield, Markus Anderljung, Miles Brundage, Julian Hazell et al., *Computing Power and the Governance of Artificial Intelligence*, CTR. FOR GOVERNANCE OF AI, https://cdn.governance.ai/Computing_Power_and_the_Governance_of_AI.pdf [https://perma.cc/D7H8-7BXP] (encouraging policymakers to narrowly tailor disclosed information).

201. *See* Bernabei et al., *supra* note 9; Onni Aarne, Tim Fist, and Caleb Withers, *Secure, Governable Chips*, CTR. FOR A NEW AM. SEC., https://www.cnas.org/publications/reports/secure-governable-chips [https://perma.cc/AUD9-ZQQE].

202. Sastry et al., *supra* note 200.