



March 3, 2023

BY ELECTRONIC SUBMISSION

Dr. Arati Prabhakar
Director, Office of Science and Technology Policy

Re: Request for Information; Digital Assets Research and Development (topics #1-6)

Dear Director Prabhakar,

Aleo Systems Inc. appreciates the opportunity to comment on the Office of Science and Technology Policy’s Request for Information on “Digital Assets Research and Development.” We are a venture-backed firm researching and developing advanced cryptographic techniques—specifically zero-knowledge proofs (“ZKPs”)—that would permit applications on distributed systems to confirm facts without unnecessarily sharing and compromising the underlying data that proves those facts.¹

This technology is critical to the future of the internet. In Part I, we explain that distributed systems—like the internet, and like blockchains—depend on the ability to exchange data confidentially, while commercial forces simultaneously work to limit anonymity. We then argue that the internet’s data security problems, which are already considerable, would be made significantly worse by the design of first-generation blockchain technology.

In Part II, we get to Aleo’s *raison d’être*: leveraging advanced cryptography to solve these problems. This technology can power the next generation of solutions for digital identity, provenance, authentication, private records, and data control in the age of an ever-sprawling “internet of things.” These solutions can fundamentally upgrade the internet.

They also fall directly within OSTP’s mandate to “kickstart research on next-generation cryptography, transaction programmability, cybersecurity, and privacy protections,”²—in addition to President Biden’s directive that the OSTP explore ways to ensure U.S. leadership in technologies of the future, especially those critical to our economic prosperity and national security.³ Having American companies lead the way means that we can build democratic values into how these systems work—precisely what happened with the internet, and precisely *why* the internet turned out to be so successful. It also means that more of the companies and jobs will be in the United States, allowing more oversight and potentially creating hundreds of billions of dollars of value. By their nature, these jobs will also be more distributed than the twentieth-century workforce, benefiting all states rather than a small handful.

¹ This comment uses the term “distributed systems” to refer to a network of unaffiliated nodes (like computers or mobile phones) that can nonetheless coordinate on activities like sending and receiving information. We use “blockchain” and “distributed ledger” interchangeably to refer to distributed (rather than centralized) ledger systems.

² *FACT SHEET: White House Releases First-Ever Comprehensive Framework for Responsible Development of Digital Assets*, THE WHITE HOUSE: STATEMENTS AND RELEASES (Sept. 16, 2022), <https://perma.cc/BQ4B-2YRG>.

³ President Joseph R. Biden, Jr., *A Letter to Dr. Eric S. Lander, the President’s Science Advisor and Director of the Office of Science and Technology Policy*, THE WHITE HOUSE: PRESS RELEASES (Jan. 15, 2021), <https://perma.cc/9MDM-FYJD>.

Part I – Privacy Enabled the Internet

The internet was a revolution in decentralized networking. It allowed individuals to connect their computers to a world wide web in order to exchange information. In its early days, many skeptics doubted the value of this technology and believed its apparent anonymity would more naturally facilitate crime and lawlessness than legitimate use cases.⁴ But subsequent history told a very different story, one that provides valuable lessons that should inform policy with respect to blockchain technology today.

First, almost all of the internet’s important use cases—email, messaging, banking, e-commerce—depend on the internet having a neutral infrastructure that facilitates the *confidential* exchange of sensitive personally identifiable information (PII) and other data. It’s hard to believe users would have adopted email or messaging applications if they believed that the entire world could read the content of their messages. Indeed, many people use web browsers simply assuming they have privacy—as they might with a doctor, therapist, or close friend, asking questions they would not ask in public.⁵

At the same time, while the TCP/IP architecture of the internet appeared neutral and anonymous—and many first understood it that way—this did not mean that all activity on the internet would be anonymous. As Larry Lessig, Tim Wu, and others have pointed out, the industries built on the internet quickly found ways to label and organize information (which in turn permitted more local regulation of the internet than people first understood).⁶ The practice of IP mapping, for instance, permits tracing IP addresses to particular geographic locations; the advertising industry, which benefits from location-specific advertising, monetized this technique soon after its discovery.⁷ Moreover, e-commerce depends on “cookies”—little bits of data that, among other things, allow website operators to know that the person “checking out” is the same person who added certain items to a shopping cart.⁸ It turned out the internet was a lot less anonymous than it looked.

In fact, the absence of data security ultimately held back innovation; this is why one of the most important policy developments for unlocking e-commerce was the government permitting the use of stronger encryption (which provides, for instance, more security when sharing payment information).⁹ The key insight is that some amount of privacy is *necessary* to realize the internet’s full possibilities, while complete anonymity limits the commercial potential of distributed technology.

There are two specific takeaways from this experience. First, concerns about anonymity—on the internet, and now on distributed ledgers—often underestimate the overwhelming commercial incentive to identify users, and thus in the long run are likely to be displaced by concerns about invasions and exploitation of privacy, as we now see with public policy concerns about the exploitation of user data.¹⁰ Second, distributed systems thrive based on finding an “enabling” amount of privacy that permits: (1)

⁴ JACK GOLDSMITH & TIM WU, WHO CONTROLS THE INTERNET xii, 13-14 (2006); *see id.* at xii (“In the 1990s, many believed that nations could not control the local effects of unwanted Internet communications that originated outside their borders, and thus could not enforce national laws related to speech, crime, copyright, and much more.”); *id.* at 3 (co-founder at MIT’s Media Lab asserting that the “internet cannot be regulated”); LAWRENCE LESSIG, CODE: VERSION 2.0, at 31 (2006).

⁵ *See, e.g.*, Michael Barbaro & Tom Zeller, Jr., *A Face Is Exposed for AOL Searcher No. 4417749*, NY TIMES (Aug. 9, 2006), <https://www.nytimes.com/2006/08/09/technology/09aol.html>.

⁶ *See* LESSIG, *supra* note 4, at 38-83; *see generally* GOLDSMITH, WHO CONTROLS THE INTERNET.

⁷ GOLDSMITH, *supra* note 4, at 7.

⁸ LESSIG, *supra* note 4, at 48-49.

⁹ *See, e.g.*, STEVEN LEVY, CRYPTO, at 312 (2002).

¹⁰ *See, e.g.*, Statement of Chair Lina M. Khan Regarding the Commercial Surveillance and Data Security Advance Notice of Proposed Rulemaking (Aug. 11, 2022), <https://perma.cc/AH3Z-SKFG>.

data security; (2) individuals to control their data and retain the dignity of choosing with whom they share this information, just as they do in the physical world; (3) commercial activities that are easy and frictionless for users; and (4) the pursuit of public policy goals like limiting abusive material, preventing illicit finance, and deterring, detecting, and punishing crime.

Finally, it is important and humbling to note that almost none of the internet's most important and useful infrastructure, much less its applications, were predictable at inception. For example, in the 1980s, McKinsey famously advised AT&T that the size of the mobile phone market in 2000 would amount to 900,000 phones, noting that such devices would be absurdly heavy and suffer low battery lives, bad service, and expensive marginal costs; of course by the year 2000, 109 million people had mobile phones,¹¹ with 900,000 joining every three days.¹² It wasn't just McKinsey. Companies as sophisticated as Kodak failed to predict the integration of mobile phones and cameras. And few if any predicted the size of markets for user-generated content until those markets came to be, though now they seem obvious—much less the ability to have mobile-powered rideshare or grocery-delivery systems. These distributed systems, once created, incentivized innovation in unforeseeable ways, leading to massive investments in software as well as the hardware that enables them.

The internet's data security problem

The problems with internet data security almost go without saying. By one estimate, approximately “4,145 publicly disclosed breaches ... exposed over 22 billion records in 2021.”¹³ In 2015, hackers stole the records of over 20 million people from the Office of Personnel Management, including fingerprint records for over 5 million people.¹⁴ In 2017, hackers accessed the personal data of 143 million U.S. consumers—roughly 44% of the U.S. population—by breaching Equifax's servers.¹⁵ IBM estimates that the average cost of a data breach is over \$9 million, with stolen credentials on average costing more and taking 327 days to identify, to say nothing of how attackers may use that information to harm victims in the future.¹⁶ These attacks amount to a serious tax on the U.S. economy.

The basic cause is the absence of data security. Consumers are often only as safe as the most negligent third parties in charge of their data. Even companies with advanced IT systems struggle to adopt basic cybersecurity measures, much less state-of-the-art practices. The overwhelming emphasis, however, has been on trying to drive the adoption of better practices. These efforts, while important, have failed to keep up. In the meantime, the distributed nature of the internet gives attackers more points of attack—a problem getting worse as more and more systems integrate to share more and more data. As a result, the expansion of cloud computing—a more distributed form of computing versus the *status quo ante* of storing and computing locally—may augur even more data breaches, because it means more data is shared. Indeed, IBM estimates that nearly half of all breaches occurred in cloud infrastructure.¹⁷

¹¹ Harry McCracken, *Shocker: In 1980, Motorola Had No Idea Where the Phone Market Would Be in 2000*, TIME MAGAZINE (Apr. 15, 2014), <https://perma.cc/R9BZ-WB3Z>.

¹² *Cutting the cord*, THE ECONOMIST: SPECIAL REPORT (Oct. 7, 1999), <https://perma.cc/B4QS-TKY9>.

¹³ *Over 22 billion records exposed in 2022*, SECURITY MAGAZINE (Feb. 10, 2022), <https://www.securitymagazine.com/articles/97046-over-22-billion-records-exposed-in-2021>.

¹⁴ See *Cybersecurity Resource Center*, OPM, <https://perma.cc/9T28-68AM>; see also Damian Paletta, *Government Personnel Cyber Breach Worse Than Previously Thought*, THE WALL STREET JOURNAL: ARTICLES (Sept. 23, 2015), at <https://www.wsj.com/articles/government-personnel-cyber-breach-worse-than-previously-thought-1443025119>.

¹⁵ Colin Dwyer, *Hackers Accessed the Personal Data Of 143 Million People, Equifax Says*, NPR (Sept. 7, 2017), <https://perma.cc/HJ8K-ZWSH>.

¹⁶ *Cost of a data breach 2022*, IBM: REPORTS, <https://perma.cc/UY6W-XSQ3>.

¹⁷ *Ibid.*

But the internet’s data privacy problems are much broader and deeper than concerns about criminal hacks. The reality is while we were initially concerned about the anonymity of the web, its infrastructure counter-intuitively supported a massive multi-hundred billion dollar data surveillance industry and the rise of surveillance-based capitalism.¹⁸ Indeed, it turned out the internet as initially conceived was *too* open, and it was the addition of *more* privacy in the form of the HTTP over SSL (“https”) that unlocked critical use cases like e-commerce.¹⁹ Even now, and even after the addition of these technologies, headlines like the Cambridge Analytica scandal²⁰ demonstrate how third parties can leverage data for purposes of fraud, scams, election interference, and social engineering to dupe the vulnerable and to turn people against their fellow citizens, neighbors, and loved ones.

The first wave of policy responses—including GDPR—represent important pioneering efforts by policymakers to at least increase awareness of the contours of the problem. But largely these have created click-through regimes that probably very few internet users read. As the current Chair of the Federal Trade Commission Lina Khan recently suggested, these procedural protections are not enough.²¹

Blockchain technology 1.0

The internet permits decentralized and unaffiliated nodes (computers, phones, devices) to connect and exchange information. Blockchain technology takes the internet one step further and permits unaffiliated nodes to cooperate to maintain and update a ledger accurately and instantaneously, even though these nodes are self-interested and otherwise have no reason to trust each other.²²

Current blockchain solutions accomplish this by essentially requiring each node to “yell the answer out loud”—in other words, all updates to the ledger are publicly announced in order to be recorded on an immutable ledger and are then accessible by anyone with an internet connection. In the context of financial applications, this is like requiring that all Venmo users use only the “public” setting (a social media feature that broadcasts the amount of the transaction and the participants to other Venmo users). The pseudonymity feature of blockchains like Bitcoin—which use your “public address” instead of your literal name or email address—provides little protection because it is a light lift for motivated observers to connect real life identities with public blockchain addresses (as firms are already doing). And once a user’s wallet address has been linked to her real identity, observers can resurrect the entirety of that person’s transactional history—a risk that does not exist in the traditional financial system. This endogenous transparency significantly limits the use cases for this technology. A financial system built on an open and transparent blockchain would be extremely concerning for individual autonomy, since purchases may say more about individuals and their identities than they would otherwise be willing (or even safe, in the case of some vulnerable groups) to share. Likewise, such a system would be unworkable for corporate data that constitutes trade secrets or that companies are not prepared to share more broadly.

¹⁸ See, e.g., *Last Week Tonight: Data Brokers*, YOUTUBE (April 11, 2022), <https://www.youtube.com/watch?v=wqn3gR1WTcA>.

¹⁹ See, e.g., *supra* note 9.

²⁰ Issie Lapowsky, *How Cambridge Analytica Sparked the Great Privacy Awakening*, WIRED (March 17, 2019), <https://perma.cc/MND2-6RAT>.

²¹ Statement of Chair Lina M. Khan, *supra* note 10 (asserting that process requirements sidestep “more fundamental questions about whether certain types of data collection and processing should be permitted in the first place”).

²² Distributed ledgers like Bitcoin and Ethereum use a native coin to incentivize nodes to do the work of verifying the accuracy of submissions to the ledger (otherwise those nodes wouldn’t waste their time and energy to do so). As a result, this means the first use case for these ledgers is keeping track of these coins. But, as many have pointed out, there is no reason why similar algorithms couldn’t use native coins to incentivize verification of non-financial applications (*i.e.*, to keep ledgers about something other Bitcoin or ETH account summaries).

The openness of these systems also limits non-financial applications. Distributed ledgers could, for example, be a more useful way to store electronic health records. In the status quo, if a healthcare patient moves or switches providers, the patient must arrange with their first doctor to transfer files to the second doctor—and risks losing medical records by failing to do so. This is because each healthcare provider functionally maintains a ledger for each patient on its own proprietary and functionally non-interoperable system (despite regulatory efforts to make these systems interoperable²³).

In other words, the status quo gives healthcare providers physical and electronic ownership of their patients' records. A distributed ledger could turn that system on its head, making patients the owners of the records, while healthcare providers could serve as nodes that update records. In that system, the records travel with the patient even if they change healthcare providers. This would be a more convenient experience for patients, one that recognizes that their healthcare data belongs to them—and that healthcare providers merely update but do not own this data. Ironically, while this “decentralizes” data (returning it to the user), techniques like homomorphic encryption, multiparty computation, and differential privacy can make that data more broadly available for scientific analysis (and even the development of AI) *without* compromising privacy.²⁴

The problem, however, is that the architecture of current blockchain technology means that each healthcare provider would “yell the answer out loud” and record that answer for all to see. Many would understandably not want this in the context of healthcare—or in many other contexts, as such records could consist of information that people or companies may not want to broadcast to others. Worse, if people adopt these systems without understanding the privacy implications—as happened with many applications on the internet—large swaths of consumer data would be available for exploitation. In that world, where it becomes too late for us to build our values into the code we use, the task for policymakers will be damage control.

Part II - Advanced Cryptography Can Provide Infrastructure-level Solutions

The data insecurity on current distributed infrastructure is staggering. The lowest-hanging fruit is information that companies store about their customers that they must store—even if they would rather not—to efficiently run their business. Take an early example: passwords. Originally many websites stored passwords on company servers. This was a risky practice because if hackers got access to those databases, they could access customer accounts. The clever and “enabling” solution was to stop storing the passwords and to instead store encrypted (“hashed”) versions of their passwords (and over time, using more sophisticated techniques like “salting” to further protect customer data against increasingly sophisticated attack strategies like dictionary attacks).²⁵ This solution, while broadly implemented for passwords, has not been implemented for the overwhelming majority of data stored on servers.

More recently, the development of advanced cryptographic technique—such as ZKPs, homomorphic encryption, multi-party computation, and differential privacy—unlock an entirely new toolkit and design space for addressing these challenges. Aleo is particularly interested in the potential of ZKPs, which allow individuals and entities to prove that something is true—with overwhelming mathematical certainty—without sharing the underlying data.

A ZKP of a particular statement has three features. First, the proof must be “complete,” meaning if someone provides the proof, we know with certainty that the particular underlying statement being

²³ See, e.g., Promoting Interoperability Programs, CMS (Jan. 9, 2023), <https://perma.cc/MBY2-ZYLE>.

²⁴ See, e.g., Understanding Differential Privacy, U.S. CENSUS BUREAU, <https://perma.cc/C46F-XAKV>.

²⁵ MIKE ROSULEK, THE JOY OF CRYPTOGRAPHY 204-205 (2021), <https://joyofcryptography.com/pdf/chap11.pdf>.

proven is true (*i.e.*, there is no further trust required). Second and related, it must be impossible for someone to provide a ZKP of a particular statement if that statement is false. Finally, the proof must be “zero knowledge,” meaning the proof should not reveal anything about the statement other than the fact that the statement is true. One classic example demonstrates how these conditions may be satisfied. Imagine you want to “prove” that you know Waldo exists on a poster, but you don’t want your counterpart (the “verifier”) to know *where* Waldo is. You could hide the poster underneath a large piece of cardboard that has a cutout that is precisely the shape of Waldo, thereby allowing you to prove Waldo exists without providing any information on his whereabouts.²⁶

Since ZKPs were first discovered in 1985,²⁷ researchers have devised techniques to reduce compute time and complexity, making ZKPs more practical to implement. Indeed, Aleo’s team and advisors have been doing research at the cutting edge of this space for many years.²⁸ As this shift happens, below are just some of the potential applications for such a technology:

- **Identity Management.** Proving identity without needing to share an actual passport or driver’s license. The process is similar to the existing workflow for creating and verifying passwords. The user takes the relevant information (*e.g.*, an email or passport) and posts it to an endpoint that creates a credential demonstrating something about that person’s identity—that it has been verified according to exacting standards, that the person is over 18, or other salient characteristics about the user. That credential can then be leveraged wherever it is accepted. This system minimizes how the underlying data (email, passport, etc.) is shared so as to significantly reduce vulnerabilities, as well as the risk that irrelevant or inappropriate factors such as race or gender are considered by the individual or entity requesting the credentials. If implemented, this technology can save millions of users from data breaches where identifying material is exposed.²⁹
- **Authentication and Provenance.** Beyond identity, this allows us to confirm the authenticity of data without necessarily stamping PII on files (*e.g.*, photographs) that are widely distributed. Computer scientists have already suggested this as a means of combating disinformation.³⁰
- **Private records that third parties can update.** *E.g.*, Health records, as discussed above.
- **Data control.** Data control as the “internet of things” blossoms. The capacity for oversharing of raw data increases exponentially as consumers use more and more devices that share information.³¹

²⁶ Another common example is using a wristband to demonstrate age (rather than sharing date of birth each time). See also *Computer Scientist Explains One Concept in 5 Levels of Difficulty*, WIRED (Jan. 1, 2022), <https://www.wired.com/video/watch/5-levels-zero-knowledge-proof>; Matthew Green, *Zero Knowledge Proofs: An illustrated primer*, CRYPTOGRAPHY ENGINEERING: BLOG (Nov. 27, 2014), <https://perma.cc/FHE7-PZ7S>.

²⁷ Shafi Goldwasser et al., *The Knowledge Complexity of Interactive Proof-Systems*, in PROVIDING SOUND FOUNDATIONS FOR CRYPTOGRAPHY: ON THE WORK OF SHAFI GOLDWASSER AND SILVIO MICALI (Oded Goldreich ed., 2019). The authors subsequently won (along with two others) the Gödel Prize for this work, a prestigious prize for outstanding papers in the area of theoretical computer science.

²⁸ See, *e.g.*, Sean Bowe et al., ZEXE: Enabling Decentralized Private Computation, in 2020 IEEE SYMPOSIUM ON SECURITY AND PRIVACY (SP) (2020), <https://ia.cr/2018/962> (multiple Aleo team members including co-founder and CTO Howard Wu demonstrating how to implement ZKPs in distributed systems).

²⁹ See, *e.g.*, Michael Rosenberg et al., zk-creds: Flexible Anonymous Credentials from zkSNARKs and Existing Identity Infrastructure, in 2023 IEEE SYMPOSIUM ON SECURITY AND PRIVACY (SP) (forthcoming May 2023), <https://eprint.iacr.org/2022/878.pdf>.

³⁰ See, *e.g.*, Trisha Datta and Dan Boneh, *Using ZK Proofs to Fight Disinformation*, MEDIUM: DAN BONEH (Sept. 29, 2022), <https://perma.cc/65X2-J289>.

³¹ See, *e.g.*, Sue Halpern, *Private Eyes*, NY BOOKS: THE NEW YORK REVIEW (March 9, 2023), <https://www.nybooks.com/articles/2023/03/09/private-eyes-the-fight-for-privacy-citron/> (explaining how Roomba’s photographs of users’ homes were lost in a data breach, including images of a customer using the toilet).

But that data can be encrypted and confirmed with ZKPs in order to protect against cyber-attacks, and yet still be made available for machine learning.

- **Proving without oversharing.** Proving things like financial health (*e.g.*, FICO scores) without needing to share bank statements or other highly sensitive specific information that can be abused, exploited, or negligently treated by the recipient.
- **Compliance tools.** ZKPs allow for greater access control and separation of responsibility in enterprise and governmental systems. They can also enable cross-validation of data to ensure accuracy and prevent fraud, trace provenance (*e.g.*, for weapons tracking), and enhance bank regulation techniques (*e.g.*, by programmatically updating new regulatory requirements or proving compliance with capital ratios and other solvency requirements in real time).
- **Secret ballot voting.** Voting in a secret ballot system where the ledger is not controlled by one entity so that third parties are more confident that the ledger has not been tampered with.
- **Digital Dollars.** The need for privacy here is apparent, and has been noted by the President’s Executive Order on Ensuring Responsible development of Digital Assets,³² the Federal Reserve,³³ the G7,³⁴ and the Digital Dollar Project.³⁵ ZKPs can allow for digital dollars in which privacy is guaranteed, while giving issuers or other intermediaries tools to comply with applicable regulations (*e.g.*, a centralized issuer can KYC users and retain decryption keys called view keys that allow compliance and audit teams to identify users and suspicious transactions).
- **Diplomacy.** In 2016, the Princeton Plasma Physics Laboratory demonstrated a technique that could allow inspectors to confirm disarmament (*e.g.*, whether an object is indeed a nuclear weapon) without recording or revealing the internal workings of the weapon, which might be secret.³⁶

The impact on semiconductors

The bipartisan CHIPS and Science Act of 2022 recognizes the national security and economic significance of the semiconductor industry.³⁷ In light of that, it’s important to understand that advanced cryptography (including ZKPs) will require computations that can be optimized at the hardware level. This is because ZKPs will require repeating similar math at the software level—and repetition at the software level benefits from optimization at the hardware level. In this case, it will soon become possible to optimize and run these operations orders of magnitude faster using specialized hardware like field programmable gate arrays (“FPGAs”) and application specific integrated circuits (“ASICs”).

This means that the country that takes the lead in leveraging advanced cryptography will also have the expertise and the incentive to optimize hardware. Aleo is on the pioneering edge of this technology, and we believe it will create an entirely new multi-billion dollar industry. We have sponsored testnets that have yielded surprising and significant levels of progress in computation times. In addition,

³² *Executive Order on Ensuring Responsible Development of Digital Assets*, THE WHITE HOUSE: PRESIDENTIAL ACTIONS (ISSUED MARCH 9, 2022), <https://perma.cc/YC66-BYLA> (mentioning privacy ten times).

³³ Federal Reserve System, *Money and Payments: The U.S. Dollar in the Age of Digital Transformation*, FEDERAL RESERVE (Jan. 2022), <https://perma.cc/D6BR-ZG8A> (“[p]rotecting consumer privacy is critical”).

³⁴ Rishi Sunak & Andrew Bailey, *Public Policy Principles for Retail Central Bank Digital Currencies (CBDCs)* (2021), <https://perma.cc/E69B-E2YK>.

³⁵ *Privacy Principles for a Digital Dollar*, DIGITAL DOLLAR PROJECT, <https://perma.cc/GAW4-JD5N>.

³⁶ John Greenwald, *PPPL and Princeton demonstrate novel technique that may have applicability to future nuclear disarmament talks*, PRINCETON PLASMA PHYSICS LAB’Y (Sept. 20, 2016), <https://perma.cc/52SF-ZQYB>.

³⁷ *FACT SHEET: CHIPS and Science Act Will Lower Costs, Create Jobs, Strengthen Supply Chains, and Counter China*, THE WHITE HOUSE: STATEMENTS AND RELEASES (Aug. 9, 2022), <https://perma.cc/9GV2-F27C>.

we helped create and sponsor Z-prize, an industry-wide effort that awarded millions of dollars to contestants who significantly improved the latency, throughput, and efficiency of computing ZKPs across multiple hardware platforms, including specialized hardware such as GPUs and FGPAs.³⁸

Aleo plans to invest significant sums of money into this technology. If—and as—applications increase exponentially, it is our view that early winners can develop a significant lead, as the markets have witnessed with Intel on CPUs or Nvidia and AMD with GPUs. In other words, this technology is not amenable to a “wait and copy” strategy; as the history of semiconductors highlights, those who take the lead in specialized technology can often retain a formidable advantage.³⁹

Support from the government

There are two areas where the government can act to help American companies innovate and take the lead in this race: inspiring demand and addressing regulatory uncertainty.

Inspiring demand. While compute times for ZKPs are significantly reducing,⁴⁰ there is still a “chicken and egg” problem with supply and demand; the supply of computing infrastructure for ZKPs is low because the demand is low. This is because there are few products to create demand, which in turn keeps the supply of computing infrastructure low. The introduction of demand signals will greatly increase investment in the space. To that end, and to stay on the cutting edge of data security, the government should explore pilot projects that leverage the technology. The Treasury Department and the Federal Reserve, in investigating a digital dollar, have already noted the importance of privacy and compliance with existing regulations.⁴¹ As mentioned above, ZKPs allow for configurable levels of privacy—*i.e.*, the government could make privacy the default but retain the ability to view transactions when authorized (*e.g.*, with a warrant, or when required by applicable law like the Bank Secrecy Act).

There are myriad potential government use cases. The Social Security Administration could launch a pilot program on digital identity. The Census Bureau already leverages differential privacy to protect the privacy of participants;⁴² it can and should explore using ZKPs to reduce data retention without losing the ability to conduct data analysis and may even be required to do so under its own disclosure avoidance regime.⁴³ The Office of Personnel Management could do the same with fingerprints and other information that it wishes to leverage but not necessarily store. These projects would spur even more innovation that could help ensure the United States is on the leading edge of applied cryptography.

Regulatory uncertainty. It’s important to understand the role of the native tokens (such as Bitcoin or ETH) in distributed ledger systems. Distributed ledgers work because their algorithm incentivizes participation by rewarding that participation—specifically the work it takes to ensure the ledger’s

³⁸ See Alex Pruden, *Announcing The Inaugural Zprize Competition Results*, ZPRIZE (Dec. 6, 2022), <https://www.zprize.io/blog/announcing-zprize-results>.

³⁹ CHRIS MILLER, CHIP WAR: THE FIGHT FOR THE WORLD’S MOST CRITICAL TECHNOLOGY (2022).

⁴⁰ Thanks in part to government efforts. See, *e.g.*, Dr. Joshua Baron, *Securing Information for Encrypted Verification and Evaluation (SIEVE)*, DARPA, <https://perma.cc/4EXH-NF3U>.

⁴¹ See, *e.g.*, *The Future of Money and Payments*, DEP’T OF THE TREASURY (Sept. 2022), <https://perma.cc/6K5T-GUC6>.

⁴² *Statistical Safeguards*, U.S. CENSUS BUREAU, https://www.census.gov/about/policies/privacy/statistical_safeguards.html.

⁴³ See 2020 Decennial Census: Processing the Count: Disclosure Avoidance Modernization, U.S. CENSUS BUREAU, <https://www.census.gov/programs-surveys/decennial-census/decade/2020/planning-management/process/disclosure-avoidance.html>.

accuracy—with native digital tokens generated by that algorithm.⁴⁴ For instance, Bitcoin pays miners with Bitcoins, and Ethereum pays its validators with ETH. Nothing about ZKPs or other privacy enhancing technologies changes the need for these tokens as part of ensuring the accuracy and security of distributed ledgers. There are significant unresolved policy questions about the regulatory status of these tokens—specifically whether they are securities and must comply with a regulatory framework designed for wholly different types of assets—as well as how to adopt risk mitigation frameworks that comply with rules designed to prevent illicit finance.

The absence of a comprehensive regulatory framework around digital tokens has created substantial uncertainty for companies wishing to innovate in this space within the bounds of the law. In this environment, the concern is that the advantage goes to unscrupulous players or those who innovate abroad. Aleo believes the United States is the best place on earth for innovation, and that good regulations can protect everyone—the public, internet users, and even the companies that those rules regulate. That’s why it’s so important for these regulations to be adapted to what is unique about distributed ledger systems, so that regulations enable responsible innovation rather than ending risk by ending innovation.

The government does not need to compromise its longstanding public policy goals to promote this kind of innovation. The primary consequence of advanced cryptography will be enabling new, better, and more secure systems. At the same time, the same bad actors who try to use the internet for scams, frauds, illicit finance, or other malfeasance will not shy away from new technologies; just as we did with the internet, we need to identify ways to cabin, deter, and punish this behavior.

Applied cryptography can help. For instance, it can help promote data security. Right now, the absence of data security infrastructure on the internet has left thousands of honeypots for hackers and cybercriminals. ZKPs can secure vast swaths of this data without compromising functionality. More fundamentally, leveraging ZKPs as authentication can help reduce identity theft and identity-related crimes, which are mushrooming as “deep fake” technology becomes more prominent. Professor Danielle Citron and others have done tremendous work highlighting the enormous scale of abuses—ranging from revenge porn to fraud—taking place online right now.⁴⁵ No doubt this is a growing problem.⁴⁶

Not all of those issues can be solved with technology, but at least some and perhaps many of them can. This is because cryptography allows us to demonstrate content is not authenticated (and thus likely faked)—or that it was shared without consent; integrating these capabilities with multiparty authentication (which can require the consent of those depicted) could create a world where platforms and providers can easily search for and remove non-consensual or faked material. These technological solutions, paired with other public policy responses, could help stem abusive practices.

Finally, it is important to understand distributed ledger technology in the context of distributed systems like the internet, even—and especially—for purposes of thinking about compliance. The notion that distributed ledgers can create an unaccountable, untraceable black box for criminal activities will—just like the notion that the internet would do the same—give way to the commercial need for identification and authentication, tools that can in turn be important instruments for law enforcement.

⁴⁴ Sina Kian, *What are cryptocurrencies good for?*, MEDIUM: SINA KIAN (Aug. 12, 2021), <https://perma.cc/DPD6-WLQN>.

⁴⁵ DANIELLE CITRON, *THE FIGHT FOR PRIVACY: PROTECTING DIGNITY, IDENTITY, AND LOVE IN THE DIGITAL AGE* (2022).

⁴⁶ Sami Quadri, *Former US ambassador says Russia is using ‘deepfakes to impersonate him’*, EVENING STANDARD (Oct. 1, 2022) <https://perma.cc/9HWP-6XCL>.

In the meantime, ample avenues exist for managing risk on distributed ledgers. First, interactions on the base layer require procuring and spending that base layer’s token (the cost of having a node process your transaction); the exchanges that allow users to purchase these coins can and should be made to KYC their customers. This alone would ensure that the overwhelming number of those holding tokens are subject to oversight—and the traces left by this activity will be difficult to obscure, especially at scale. Second, these systems all depend on nodes and validators for their security; insofar as these validators can be encouraged to domicile in the United States, they can play an important part in setting standards and preventing bad behavior. For instance, following a large theft of ETH tokens, many validators collectively decided to fork the blockchain and create a ledger in which the theft had never occurred, thus effectuating a technological version of “restitution” (restoration of funds to victims).⁴⁷ Third, the authentication potential of ZKPs can help upgrade outdated KYC practices and provide a greater check on bad actors. They can be updated and refreshed more frequently and, even with ZKPs, identity solutions can build in mechanisms for ensuring compliance.⁴⁸

Finally, regulators can continue to achieve compliance at the application level.⁴⁹ Over the past ten years, public discussion has focused on the infrastructure layer because these systems—like Bitcoin and Ethereum—were new and exciting, and because many speculated on their value by purchasing their local coins. But the future of distributed systems depends not on their infrastructure, but on the applications their infrastructure enables. These applications will overwhelmingly be subject to regulations in their current form—and where they are not, it is incumbent on policymakers and the private sector to work together to enable new paradigms while addressing important public policy goals.

The history of distributed systems—whether telegram or the internet—has demonstrated that these systems are always more regulable than we think.⁵⁰ The best way to address the challenges of new technology is to invest in understanding them, so that we can understand how to achieve our public policy goals. Aleo is betting on the future of advanced cryptography, and betting that America will be and should be its home. We welcome any conversations on the best way to make that happen.

Respectfully submitted,

Sina Kian

Sina Kian

Chief Operating Officer & General Counsel, Aleo Systems, Inc.

Cc: Counsel to Aleo on this matter:
Wilmer Cutler Pickering Hale and Dorr LLP
Zachary Goldman, Esq.
Tiffany Smith, Esq.
Jason Raymond, Esq.

⁴⁷ CAMILA RUSSO, *THE INFINITE MACHINE*, ch. 21 (2020).

⁴⁸ *See, e.g.*, LESSIG, *supra* note 4, at 70 (conceptualizing identity solution that “could radically increase privacy, as well as security, for all except those whose behavior can legitimately be tracked”).

⁴⁹ *See generally* Miles Jennings, *Regulate Web3 Apps, Not Protocols*, A16ZCRYPTO (Sept. 29, 2022), <https://a16zcrypto.com/web3-regulation-apps-not-protocols/>.

⁵⁰ *See, e.g.*, GOLDSMITH, *supra* note 4, at 124.