



**LF** DECENTRALIZED TRUST

---

# Zero Knowledge Proofs for the Decentralized Trust Graph

Sanjam Garg (UC Berkeley) & Hart Montgomery (Linux Foundation)

## Graphs and Things to Prove on Them

Personhood credentials: Cryptographers  
Linux kernel web of trust

## The Solution: SNARKs Tools for Building

zkbk  
Other relevant LFDT labs

# CONTENTS

# Personhood credentials: Artificial intelligence and the value of privacy-preserving tools to distinguish who is real online

Steven Adler,<sup>\*†1</sup> Zoë Hitzig,<sup>\*†1,2</sup> Shrey Jain,<sup>\*†3</sup> Catherine Brewer,<sup>\*4</sup> Wayne Chang,<sup>\*5</sup> Renée DiResta,<sup>\*25</sup> Eddy Lazzarin,<sup>\*6</sup>  
Sean McGregor,<sup>\*7</sup> Wendy Seltzer,<sup>\*8</sup> Divya Siddarth,<sup>\*9</sup> Nouran Soliman,<sup>\*10</sup> Tobin South,<sup>\*10</sup> Connor Spelliscy,<sup>\*11</sup>  
Manu Sporny,<sup>\*12</sup> Varya Srivastava,<sup>\*4</sup> John Bailey,<sup>13</sup> Brian Christian,<sup>4</sup> Andrew Critch,<sup>14</sup> Ronnie Falcon,<sup>15</sup> Heather Flanagan,<sup>25</sup>  
Kim Hamilton Duffy,<sup>16</sup> Eric Ho,<sup>17</sup> Claire R. Leibowicz,<sup>18</sup> Srikanth Nadhamuni,<sup>19</sup> Alan Z. Rozenshtein,<sup>20</sup>  
David Schnurr,<sup>1</sup> Evan Shapiro,<sup>21</sup> Lacey Strahm,<sup>15</sup> Andrew Trask,<sup>4,15</sup> Zoe Weinberg,<sup>22</sup> Cedric Whitney,<sup>23</sup> Tom Zick<sup>24</sup>

<sup>1</sup>OpenAI, <sup>2</sup>Harvard Society of Fellows, <sup>3</sup>Microsoft, <sup>4</sup>University of Oxford, <sup>5</sup>SpruceID, <sup>6</sup>a16z crypto,  
<sup>7</sup>UL Research Institutes, <sup>8</sup>Tucows, <sup>9</sup>Collective Intelligence Project, <sup>10</sup>Massachusetts Institute of Technology,  
<sup>11</sup>Decentralization Research Center, <sup>12</sup>Digital Bazaar, <sup>13</sup>American Enterprise Institute,  
<sup>14</sup>Center for Human-Compatible AI, University of California, Berkeley, <sup>15</sup>OpenMined,  
<sup>16</sup>Decentralized Identity Foundation, <sup>17</sup>Goodfire, <sup>18</sup>Partnership on AI, <sup>19</sup>eGovernments Foundation,  
<sup>20</sup>University of Minnesota Law School, <sup>21</sup>Mina Foundation, <sup>22</sup>ex/ante, <sup>23</sup>School of Information, University of California, Berkeley,  
<sup>24</sup>Berkman Klein Center for Internet & Society, Harvard University, <sup>25</sup>Independent Researcher

August 2024

# Web of Trust: Cryptographers



$\sigma(\text{Dan}, \text{"Sanjam is a cryptographer"})$




$\sigma(\text{Sanjam}, \text{"Dan is a cryptographer"})$





"I want to prove that I'm a well-endorsed cryptographer."

"I don't want to reveal who has vouched for me as a cryptographer."





# Web of Trust: Cryptographers



$\sigma(\text{Dan, "Sanjam is a cryptographer"})$




$\sigma(\text{Shafi, "Sanjam is a cryptographer"})$



$\sigma(\text{Brent, "Sanjam is a cryptographer"})$



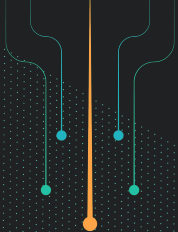
$\sigma(\text{Abhishek, "Sanjam is cryptographer"})$



$\pi :=$  "I have a valid root credential and  $\geq 4$  others with valid root credentials have endorsed me as a cryptographer"







# Can't do with Traditional Anonymous Credentials!




# XZ Utils Attack

## From Wired:

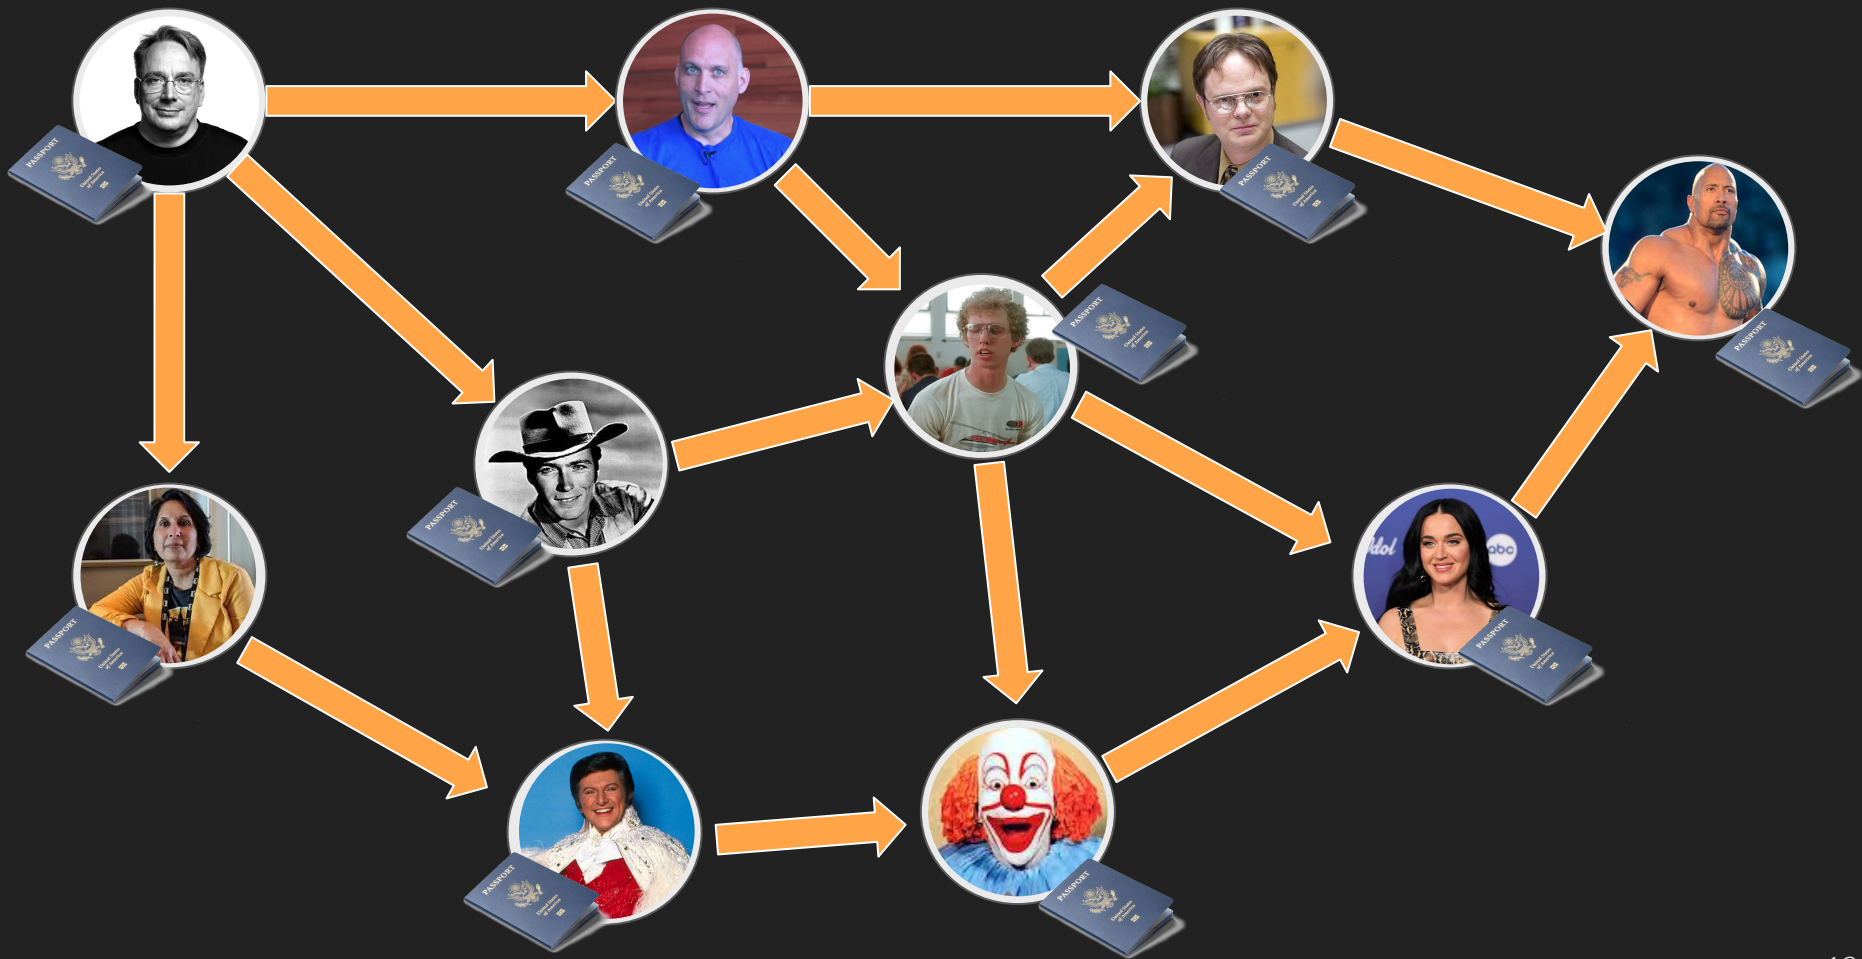
On Friday, a lone Microsoft developer rocked the world when he revealed **a backdoor had been intentionally planted in XZ Utils**, an open source data compression utility available on almost all installations of Linux and other Unix-like operating systems. The person or people behind this project likely spent years on it. They were likely very close to seeing the backdoor update merged into Debian and Red Hat, the two biggest distributions of Linux, when an eagle-eyed software developer spotted something fishy.

"This might be the best executed supply chain attack we've seen described in the open, and it's a nightmare scenario: malicious, competent, authorized upstream in a widely used library," software and cryptography engineer Filippo Valsorda said of the effort, which came frightfully close to succeeding.

# Web of Trust: Linux Kernel Maintainers



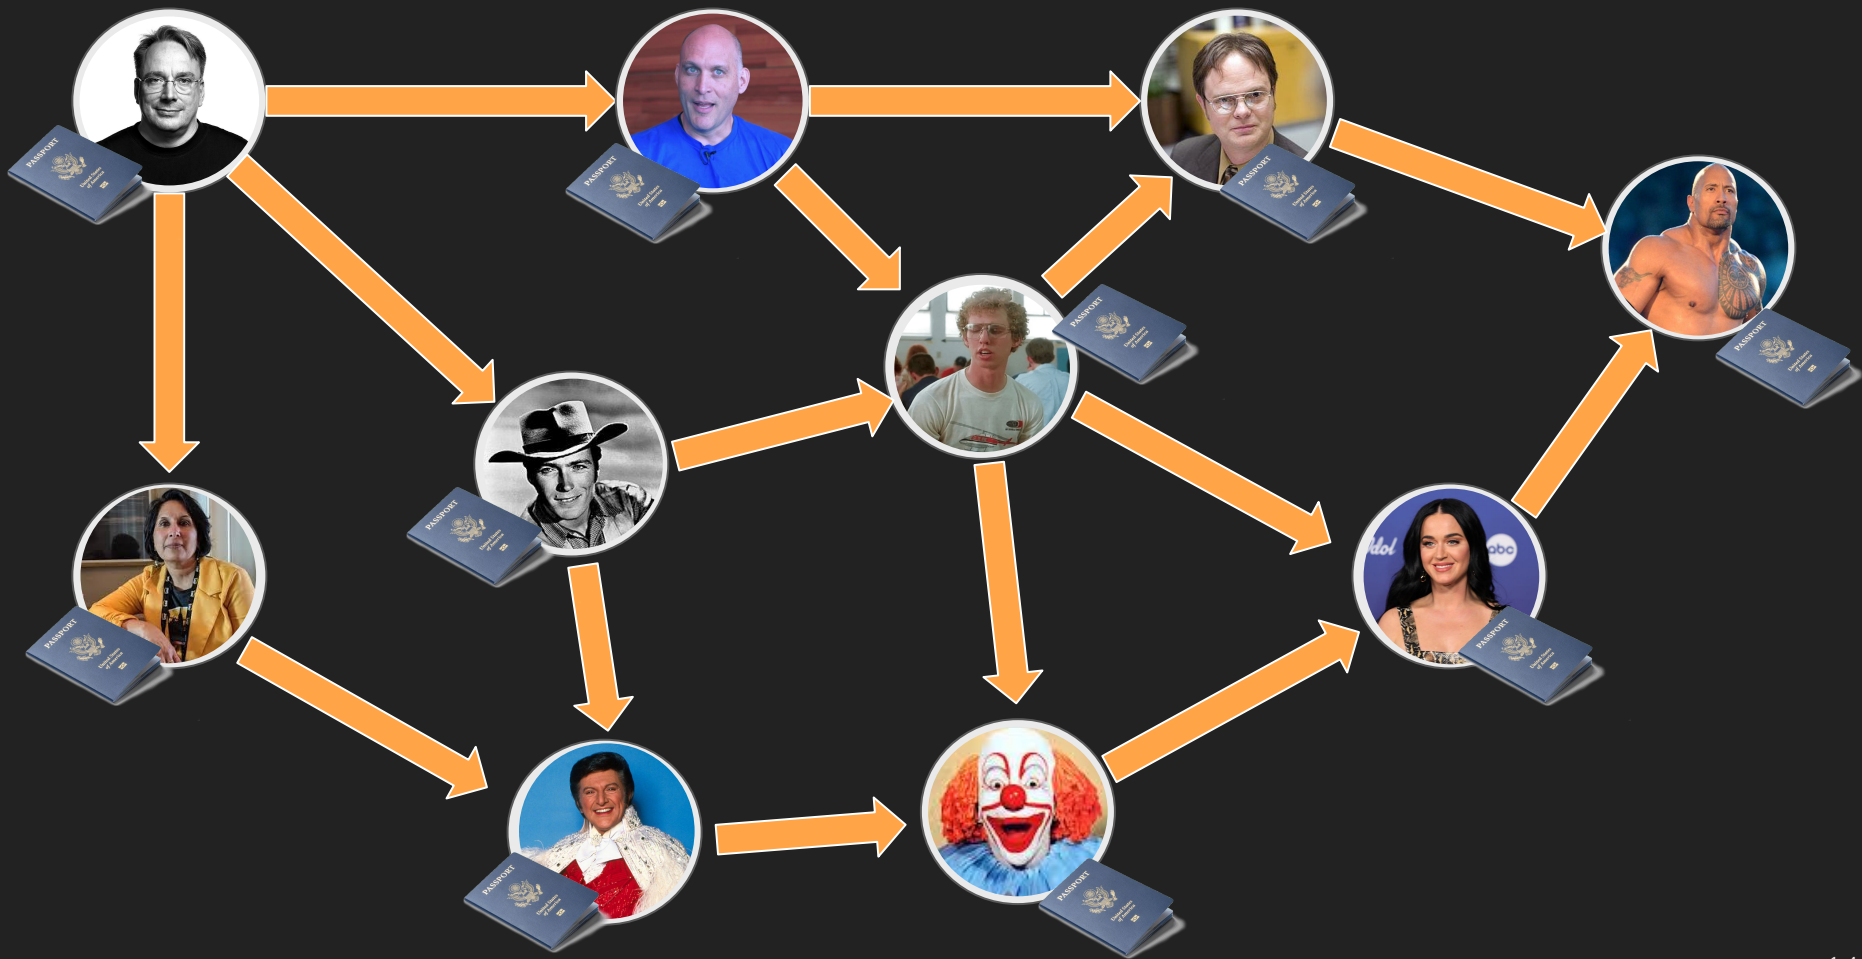
$\sigma(\text{Linus}, \text{"I know Greg as a developer and have met him in person"})$






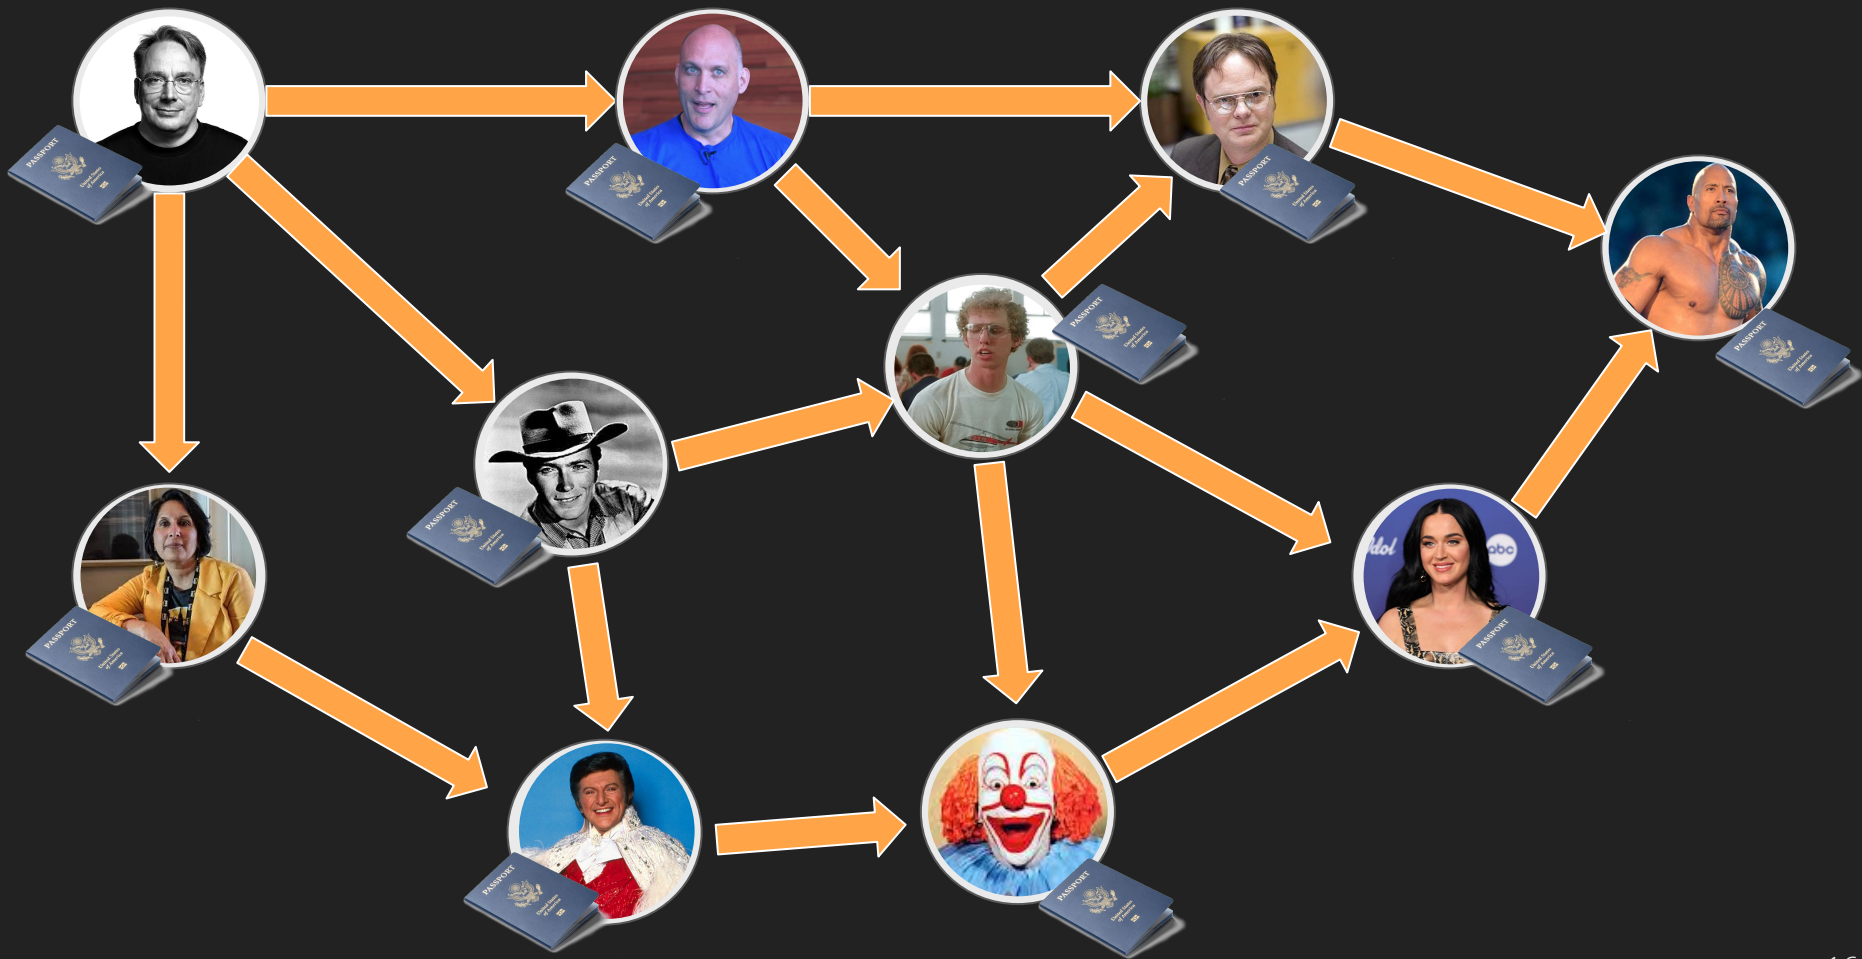
"I want to prove that I have two social paths of length at most 5 to Linus where all users on the path also have this property."

"I don't want to reveal these paths because it may open the kernel up to social engineering attacks."











# Web of Trust: Linux Kernel Maintainers




$\sigma$ (Dwight, "I know the Rock as a developer and have met him in person")

$\pi$ (Dwight, "I have an appropriate path and credentials")\*

$\sigma$ (Katy, "I know the Rock as a developer and have met him in person")

$\pi$ (Katy, "I have an appropriate path and credentials")\*



$\pi$  := "I have two valid paths to Linus with  $\leq 5$  length where all users on the path have valid credentials and two valid paths to Linus."



Easy mode: ignoring **disjoint paths**  
and other properties for now.




# ALSO can't do with Traditional Anonymous Credentials!



# SNARKs to the Rescue!

# Zero Knowledge Proof System



► **Syntax:** Two algorithms,  $P(1^n, x, w)$  and  $V(1^n, x)$ .

► **Completeness:** Honest prover convinces an honest verifier with *overwhelming* probability.

$$\Pr[V \text{ outputs } 1 \text{ in the interaction } P(1^n, x, w) \leftrightarrow V(1^n, x)] = 1 - \text{neg}(n)$$


► **Soundness:** A PPT cheating prover  $P^*$  cannot make a Verifier accept a false statement.  
For all PPT  $P^*$ ,  $x$  such that  $\forall w, C(x, w) = 0$  then we have that

$$\Pr[V \text{ outputs } 1 \text{ in the interaction } P^*(1^n, x) \leftrightarrow V(1^n, x)] = \text{neg}(n)$$

► **Zero-Knowledge:** The proof doesn't leak any information about the witness  $w$ .  $\exists$  a PPT simulator  $S$  that for all PPT  $V^*$ ,  $x, w$  such that  $C(x, w) = 1$ , we have that  $\forall$  PPT  $D$ :

$$|\Pr[D(V^* \text{'s view in } P(1^n, x, w) \leftrightarrow V^*(1^n, x)) = 1] - \Pr[D(SV^*(1^n, x)) = 1]| \leq \text{neg}(n)$$

# Succinct Non-Interactive Argument System (SNARG)




- **Completeness:** An honest prover should be able to convince an honest verifier with *overwhelming* probability.
- **Soundness:** A PPT cheating prover cannot generate an accepting proof for a false statement.
- **Zero-Knowledge:** The proof doesn't leak any information about the witness  $w$ .
  - Not all applications need zero knowledge, e.g. zk-rollups.

# Commonly Used Proof System Frameworks

Framework	Arithmetization	Algorithm	Field	Other Configs
Circom + snarkjs / rapidsnark	R1CS	Groth16	BN254 scalar	
gnark	R1CS	Groth16	BN254 scalar	
Arkworks	R1CS	Groth16	BN254 scalar	
Halo2 (KZG)	Plonkish	KZG	BN254 scalar	
Plonky2	Plonk	FRI	Goldilocks	blowup factor = 8 proof of work bits = 16 query rounds = 28 num_of_wires = 60 num_routed_wires = 60
Starky	AIR	FRI	Goldilocks	blowup factor = 2 proof of work bits = 10 query rounds = 90

From the Ethereum Foundation





# Code You Can Use




zkbk is a collection of zero-knowledge gadgets built on top of the arkworks ecosystem to facilitate **first person credentials**.

They will enable us to:

- Put privacy-preserving digital credentials in the hands of every Internet user.
- Create ubiquitous secure personal private channels.
- Facilitate strong, confidential trust relationships without intermediaries.
- Empower all of us with personal AI agents we can trust.

And we believe that doing all of this—without surveilling us or stealing our data—can make a real difference.

“zkbk” for “**Z**ero **K**nowledge **B**er**K**eley”: led by the cryptography group at Berkeley (Sanjam Garg’s group) and some of their alums.




## Other Exciting New Labs



# Generic zk-SNARKS

This is an effort led by the **Privacy & Scaling Explorations (Ethereum Foundation) zkID team** to standardise generic zk-SNARKs.

**Goal**: standardize the high-level components of a generic zk-SNARK, namely: arithmetisation, polynomial interactive oracle proof, Fiat-Shamir transform, and polynomial commitment scheme. The specification will also include information about the secure composition of these components, in particular with regard to their soundness and zero-knowledge guarantees. The aim is not to enshrine a normative proof system, but rather illustrate the usefulness of the standard by applying it to a few popular existing proof systems.




# Nightstream

NightStream is an open-source initiative developing a **post-quantum zero-knowledge proof system** that balances performance, security, and practical deployment constraints. The lab focuses on creating a **transparent, lattice-based architecture** that provides realistic performance improvements while maintaining post-quantum security guarantees.


Initial committers include people from Input Output, Stanford, Midnight, and Shielded.

- <https://github.com/dabo> (!!)
- <https://github.com/solegga> (!!)



# Nightstream Scope

- Core Research and Development:
  - Development of a lattice-based folding protocol (NEO, LatticeFold+ or similar) for post-quantum security
  - Incorporation of lookup-centric arithmetization for computational efficiency
  - Creation of a unified compilation target supporting R1CS, PLONKish, and AIR circuits
  - Research into practical aggregation techniques for multiple proofs
- Performance Optimization:
  - Hardware-software co-design for CPU and GPU acceleration
  - Small-field arithmetic optimization using the 64-bit Goldilocks field
  - Memory-efficient proving strategies for large-scale computations
  - Incrementally Verifiable Computation (IVC) support for long-running processes
- Developer Experience:
  - Rust-native SDK with comprehensive tooling
  - Migration paths from existing proof systems (including Halo 2, PLONK, STARKs)
  - Debugging and development tools for circuit construction
  - Documentation and learning resources for the ZK community
- Practical Deployment:
  - Multi-platform verification support (EVM, Move, WASM, Plutus, Compact (AKA Project Mitra))
  - On-chain integration strategies with realistic gas cost targets
  - Proof aggregation for blockchain scalability applications
  - Data availability proof generation for rollup systems
- Security and Standards:
  - Post-quantum cryptographic parameter selection and validation
  - Formal security analysis of critical components and protocols
  - Community-driven security audits and peer review
  - Compliance with emerging post-quantum standards





# Any Questions?

[sanjamg@berkeley.edu](mailto:sanjamg@berkeley.edu)

(Join [bayareacryptoday](#), our [online seminars](#))

[hmontgomery@linuxfoundation.org](mailto:hmontgomery@linuxfoundation.org)

