WORLDCOIN

Private by Design

Contents

Ecosystem

The Worldcoin project consists of multiple actors and tools, which combine to form a human-first identity network that enables trust when transacting or communicating online.

Worldcoin is a project that encompasses World ID, an anonymous digital passport, and a network that enables the use of digital assets, providing inclusive access to the global digital economy for billions of people.

> **World ID** is a decentralized identity protocol for proving you are a unique person. You can use World ID to prove you're a human in any online activity, such as authenticating videos and protecting against deepfakes. You can also use it to sign in to websites and applications, similar to Sign In with Google, proving you're a unique human without ever having to share personal data like your name, email or phone number.
>
> **World Chain** is a soon-to-be-launched layer 2 rollup on the Ethereum network that leverages World ID to prioritize human-centric transactions over bots.
>
> **WLD** is the Worldcoin token, which is freely given to individuals for being human and a part of the Worldcoin network.[1]

Worldcoin Foundation is a nonprofit organization that serves as the steward of the Worldcoin protocol. It also owns and governs most assets related to the Worldcoin brand, including intellectual property for the orb and the protocol's open-source technology.

Tools for Humanity (TFH) is a technology company that makes tools for Worldcoin, including the orb and World App.

> An **orb** is a special camera that verifies unique personhood and provides you with the data to confirm your personhood with the use of a World ID.
>
> **World App** is a Worldcoin self-custodial wallet that provides a home for World ID. With the app, you can also send and receive Worldcoin tokens and other digital funds.

For more information: What is Worldcoin, and how does it work?

[1] In eligible jurisdictions

Privacy in the Age of AI

A 2022 report by the European Union's law enforcement agency Europol suggested that as much as 90% of the internet's content could be synthetically generated by 2026.

Catfishing. Scambots. Deepfakes. Identity theft. Disinformation. The internet can be a dangerous place. The advancement of AI will make the internet more useful than ever but we must be clear-eyed about its potential to amplify existing problems. We've seen what can go wrong when a human pretends to be someone else. What happens when we think we're dealing with a human that is actually an AI agent?

What we need is a way to verify that the people we talk to, send money to, and view content from online (to name but a few examples) are actually people. Worldcoin aims to give you agency over each of these elements in the age of AI.

Preventing AI from flooding the internet with fake people requires extraordinary security. Yet many of the envisioned approaches to dealing with this are heavy-handed. There is a temptation to reach for the same old tools – removing privacy and relying on identification or verification techniques that can be co-opted for corporate and government surveillance. A panopticon might work but at what cost?

We believe there's a better way: The Worldcoin project is about creating secure technologies that put humans at the center, allowing them to better control and trust their online experiences without sacrificing their privacy.

Worldcoin is **private by design**.

How AI poses risks to privacy

We've already seen deepfakes circulate of politicians and celebrity personalities—people for whom there's a large well of online content to draw from.

But advanced AI means that deepfakes—extremely realistic video and audio—could soon be used to imitate everyday people. We're especially concerned that the internet and its users aren't prepared for the challenges advanced AI will pose. Imagine being on a Zoom call with colleagues who aren't actually colleagues or loved ones who aren't people at all. People can easily be scammed into sending money, divulging secrets, or things we have yet to imagine.

To combat malicious applications of AI, platforms must be able to know someone is a unique human. Tools such as "CAPTCHA" are no longer effective and rely on data tied to a person's digital footprint. World ID is a privacy-preserving alternative to digital interactions such as CAPTCHAs. It also serves as an alternative to KYC and digital ID systems that either reveal an individual's identity or link a person's identity to their digital activity.

Worldcoin Privacy Principles

The Worldcoin community is building something unprecedented: a globally trusted, maximally inclusive, privacy-preserving network for proving personhood. But the project's approach to privacy is counterintuitive and novel, presenting a fundamentally new route that no company or organization, or governments, has taken:

Worldcoin does not want to know who you are, just that you are a unique human.

AI-powered bots are becoming more ubiquitous, increasing in sophistication, undermining trust online and impersonating people. So knowing if you're interacting with a human or bot is increasingly important. If proving personhood online were our only concern, the solution might appear fairly straightforward: use government-issued IDs to verify our identities and navigate online. After all, we're often asked to show ID at a bank. Aren't online transactions in the age of AI potentially just as sensitive?

Leaving aside the fact that an estimated 850 million people don't have any form of official ID and that governments themselves have deployed sophisticated disinformation campaigns online, verification via government-issued IDs online is not the solution. It reveals much more about you than needed, such as your address, and puts you at risk of having your identity stolen or used for nefarious purposes, including in ways we can't imagine because of AI.

We need proof of personhood *and* privacy in the same package (with as few barriers to entry as possible).

"Until now, anyone wishing to prove their humanness online has used means such as government IDs, which are burdened by the disadvantage of identifying the user and of revealing a large quantity of other personal data, even though this is not necessary. In contrast, World ID allows an anonymous 'proof of unique humanness' and thus counters the model of 'surveillance capitalism' with a model that promotes data protection. World ID hereby strengthens the opportunities for data protection-compliant online activities."

Dr. Stefan Brink, former German State Commissioner for Data Protection and Freedom of Information in Baden-Württemberg from January 2017 to December 2022.

This is the challenge Worldcoin seeks to solve, and it does so primarily through World ID. World ID is a global digital passport that lives locally on its holder's smartphone and allows someone to prove they are a unique person without sharing personal data with anyone.

It's a humanness-verification system for the internet that allows people to be anonymous wherever they go. A verified World ID does not collect or link to identifiers like name or email, link to wallet transaction data or reveal whose World ID is used. By design, Worldcoin is founded on the practice of data minimization. It does not store identifying information.

Moreover, World ID is built so that people can use it across applications without those apps tracking their activity from one to another. There's no central repository of usage history. You can use World ID on hundreds of different apps without one ever knowing about the others —or World ID knowing about those apps.

Worldcoin is private by design with four interwoven privacy principles:

Principle 1

Security: Secured by math

Principle 2

Anonymity: Move freely online

Principle 3

Choice & Control: Your data, your rules

Principle 4

Transparency: Built in the open

Principle 1

Security: Secured by math

Without security, there is no privacy.

Worldcoin is about enabling people to be online without having their identities exposed and empowering them to distinguish between bot-based interactions and human ones. Security helps make sure that level of privacy is achieved—every time, without fail.

World ID uses many security techniques to ensure the data security of World ID holders.

One set involves human tools, such as open-sourcing and audits (see: Transparency), which help validate and pressure test the security measures that have been created and implemented as part of the Worldcoin project.

The other set includes cryptographic tools, such as ZKPs and SMPC (see: Anonymity), which use advanced mathematics to secure data, encrypt it, and keep it private or render it anonymous.

SMPC is one of the few results in cryptography that can provide perfect secrecy. ZKPs, meanwhile, use nullifier hashes, or unique values, for each application so that people's usage history can't be tracked.

Call it **secured by math**.

Principle 2

?

But using websites and apps shouldn't involve giving away our data any more than cooking dinner does. People should be able to move freely online.

Anonymity: Move freely online

Much of the time, we don't need identification just to be in the world. Cooking dinner, reading a book, sleeping. Performing these activities rarely requires us to prove who we are. For the majority of our lives, our actions go unnoticed, unobserved, and unrecorded. We are anonymous.

Anonymity is harder to attain online. Sites can see our activity, and browsers can track our online movements and behavior. This surveillance becomes potentially more harmful when we're asked to prove who we are, all the way from having our IP address monitored at the low end to displaying government-issued ID authentication at the extreme end.

But using websites and apps shouldn't involve giving away our data any more than cooking dinner does. People should be able to **move freely online**. Signing into a site with a World ID allows this.

To maintain anonymity in an online world, Worldcoin uses many privacy-preserving technologies including secure multiparty computation (SMPC) and zero-knowledge proofs (ZKPs).

Secure multiparty computation (SMPC)

It only takes a smartphone to create a World ID, but proving the holder is a unique human who hasn't created multiple World IDs—while keeping their identity private is a complicated challenge.

Biometric data, when appropriately anonymized, provides the solution. The very usefulness of biometric data means it should be collected and used minimally and when it is the only viable path, it must be handled with care.

Worldcoin does this via SMPC.

When a person verifies their World ID via an orb, the orb takes pictures of their iris and face. It uses these pictures to make an iris code, which is essentially a series of 1s and 0s. No two iris codes are the same nor do they reveal direct identifiers such as name, gender, age, etc.

This iris code is split into different pieces and permanently encrypted using SMPC, which makes data anonymous by splitting it into multiple abstracted values (SMPC shares) and storing them in separate locations managed by two legally distinct entities. In the near future additional storage partners (including universities and non-profits) will be added, meaning that the iris codes will be split into even more abstracted values stored and managed by even more independent entities. No single party has access to a part of an iris code. Rather they only have access to the SMPC share stored under their control and management.

While storing the data in multiple locations would seemingly increase the likelihood of that data being stolen, the opposite is true. The SMPC shares are stored in a way that, were a malicious actor to somehow get access to one SMPC share, it would be indecipherable; they only make sense when all of the pieces are put together.

Why store the encrypted SMPC shares at all? So that the Worldcoin protocol can keep proving a person is unique. Without it, users would need to re-verify their World ID every time an app called for it.

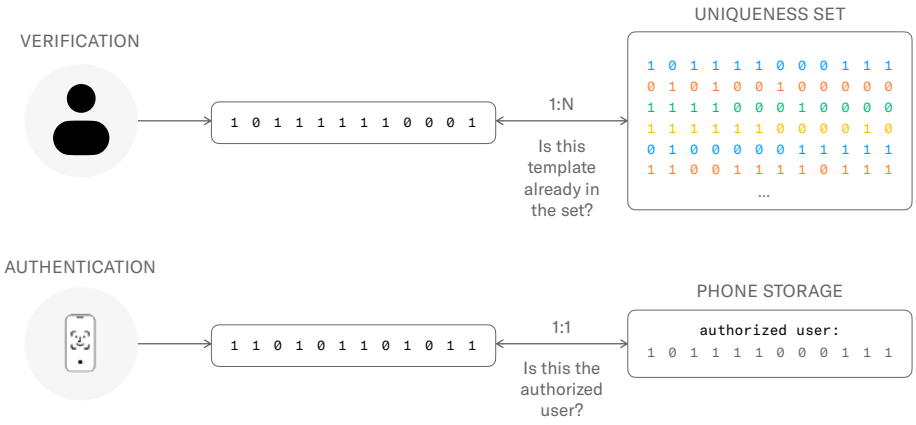

The photos themselves don't stay on an orb either. Instead, the orb encrypts the data end to end with a public key provided by the user's smartphone (no one but the user has the private key to decrypt) and then the orb transmits the encrypted pictures to the user's device before deleting them from the orb. All of this happens in a matter of seconds during the verification process.

Why irises?

Orbs—advanced cameras—are the first hardware devices to support the Worldcoin protocol. For now, the only method for verifying World IDs as belonging to a unique human is to visit an orb and take a picture of your eyes.

TFH studied many different types of biometrics for verifying World IDs, each with its own pros and cons. To meet the future proof criteria needed to provide humanness for all individuals on earth in an AI world, any method must be 1) accurate, 2) privacy-preserving, 3) extremely difficult to spoof, 4) scalable, and 5) very easy to use.

Fingerprints are very usable but easy to spoof. Irises, however, are accurate, usable, scalable, broadly inclusive, and extremely difficult to fake. And, they are privacy-preserving. There is no large public record of irises (like faces on social media), one cannot photograph a closeup of someone's iris without that person noticing, and specialized camera equipment is required to even attempt to do so.

The Worldcoin protocol is open and decentralized, and additional verification mechanisms will only strengthen the appeal and security of the project.

Zero-knowledge proofs (ZKPs)

Once a person has a verified World ID, they can use it to log in to and transact with third-party apps that integrate with the World ID protocol.

But that doesn't mean people share their World ID with the third-party app.

Instead, the World ID is used to create a disposable version of itself, similar to Apple's Hide My Email or virtual credit cards. Imagine having a company credit card that can only be used for one purchase. To pay for things at other vendors, you'd need to generate separate cards. While this might sound onerous, the protocol does it quickly and seamlessly behind the scenes, resulting in a system that is secure for apps and protects users.

The method for doing this—so that neither Worldcoin nor applications can track usage history or across apps or the parties involved in a given transaction—is a ZKP, a cryptographic tool that allows someone to prove something is true without revealing any of the information used to come to that conclusion.

Anytime you use a third-party app, the app will ask for a proof from your device. Think of it as asking your device a question. It wants to know, does this device control this World ID? The World App on the user's device sends back a ZKP that demonstrates the World ID is verified.

ZKPs go well beyond any jurisdiction's regulatory requirements. Worldcoin implemented them because they're the best way to ensure users can remain anonymous (unless the user provides additional information directly to the third party service requesting the proof) and apps can't track them. They prevent third parties—and Worldcoin Foundation itself—from ever knowing a user's World ID, or what services they interact with.

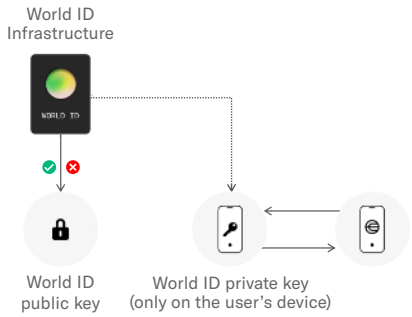

A practical analogy for ZKPs

ZKPs are advanced cryptography, so they don't lend themselves to perfect analogies. But to get the basic idea, imagine a puzzle in which you have to find a particular object or person inside a scene.[2]

After searching for a few minutes, one person says to the other that she knows where the item is. The second person, who doesn't know where the item is, isn't sure whether to believe this.

So the first person says she can prove it—without just revealing where the item is.

She photocopies the complete puzzle, cuts out the item, and holds it up to the second person, who is now certain the item is in the puzzle and that the first person knows where it is.

In this scenario, the second person is any app trying to confirm that a user is human; the cut-out item is the ZKP; and the first person is the Worldcoin protocol proving that the user is human without revealing their identity.

[2] E.g., "Where's Waldo/Wally?"

Use case:
X bot or human?

Elon Musk purchased X (then called Twitter) promising to drive bots from the platform. But the new owner quickly realized this was easier said than done. “It is extremely difficult to stop bots without affecting real users,” Musk wrote in December 2023. “As advanced AI becomes available to anyone, it will become almost impossible.”

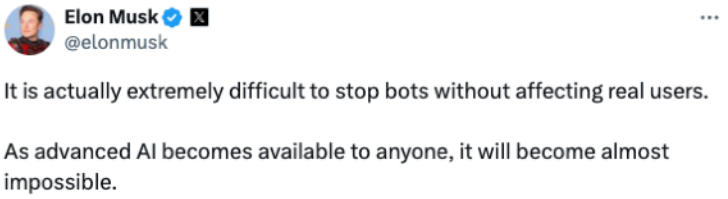

Musk was right about the challenge bots presented to the social platform. but there *is* now a viable solution for combating them without affecting users.

X allows three sign-in modes: Sign in with Google, Sign in with Apple, or a username/password. To create an account, users are asked to provide their name, phone number or email, and date of birth. Since it’s relatively easy to create multiple email addresses, it’s also quite simple to make multiple X accounts. Thus, there’s not much of a barrier for bots to get online, where they can degrade the user experience for humans.

However, World ID does not rely on email addresses, which are easy to set up, or phone numbers, which are simple to spoof. Only humans can get a verified World ID—and, unlike emails, they are limited to one. Thus, if X were to use World ID as its personhood verification mechanism, the service could add a badge that denotes such accounts were verified humans. Any bots that logged on would be unable to verify.

Importantly, doing so would even enhance anonymity for X users, who would not have to provide any additional information beyond what they have already provided to X.

This isn’t entirely theoretical. TFH built a World ID Telegram integration to get rid of spam bots on the network. Public chat administrators can mandate that individual accounts must first verify with World ID before posting in a group.

Principle 3

Choice & Control: Your data, your rules

People have become increasingly accustomed to a Big Tech paradigm in which, in exchange for accessing services, they surrender their personal data to corporations so that it can be sold to the highest bidder.[3]

Worldcoin operates outside of this paradigm. This isn't just a promise. Worldcoin is being intentionally designed to make it impossible to do so.

The Worldcoin approach is: **Your data, your rules.**

Data minimal

The starting point for recognizing people's control over their data is not asking for much data to begin with. Because verified World IDs are anonymous, people do not provide their name, phone numbers, address, or other information commonly captured by technology companies. Think about it: To get a *library card*, people need proof of residency with an address. To get a verified World ID, a global digital passport, they just need a smartphone and to visit an orb.

Use case:
Shopify discount codes

Shopify is an e-commerce platform for businesses looking to grow sales and manage online payments. Merchants on the platform sometimes look to attract new customers by issuing one-time discounts.

The problem is that people can game the system by making fake emails and claiming discounts multiple times—or deploying bots to do this for them. Instead of attracting new customers, the vendors are subsidizing a scam.

By integrating World ID in their store, merchants can let real customers scan a QR code that verifies their World ID and applies a discount code. This method ensures one discount per human, solving the merchants' problem without asking the user to provide any additional information. (Although they will have to give their credit card and shipping details to check out!)

[3] Literally! For more information on this, see https://en.wikipedia.org/wiki/Real-time_bidding

Personal Custody

The process of verifying a World ID requires the use of *some* data, namely, biometric images and iris codes. The iris codes, after going through SMPC, go to servers where they reside in a fully anonymized format. But the data provided by the user to verify personhood via the orb is not retained or provided to any third party. Instead it lives only on the person's smartphone, where it is encrypted with an individual's public key.

With Worldcoin Personal Custody, people *control* the data collected and generated during verification—including the World ID and the images—and decide who to share it with.

Face Auth

By instituting World ID, platforms can guard against bots without invading the privacy of their customers. A verified World ID provides these platforms with a high level of confidence that a user is actually a person.

But in some high-stakes scenarios (e.g., financial transactions) when platforms or people must know not only that they are dealing with a person but that they're dealing with a *specific* person. They want to know that the person using the World ID is the *same unique human* that verified the World ID on that device.

To accomplish this, apps on Worldcoin can use Face Auth.

Face Auth is a method for comparing the image taken during verification with an image of the person seeking to use the World ID and is device agnostic.

The first image is generated when a person verifies their World ID at an orb. The user's phone is by default the only place where this data exists. Those high-resolution photos are encrypted, and securely sent to the user's phone as part of Personal Custody, and completely deleted from the orb.

The second image is a selfie taken on the user's device within the World App when the user seeks to access or use their World ID. Face Auth for World ID compares the image taken by the user's device to the original face authentication image taken during verification at the orb. The user can only continue with their login or transaction if the two images match.

This prevents fraud by defending against a malicious actor who steals (or purchases) someone's phone and uses their World ID. With Face Auth, the intended user is always in control of their data and their World ID. **The comparison is done locally on the individual's device**. As a result, neither the selfie, the picture from the orb, nor any other personal data is shared with any third parties including Tools for Humanity or the Worldcoin Foundation.[4]

[4] In the future, the project will allow for people to opt-in to voluntarily share their information with Worldcoin to assist with security and AI-training. This is 100% optional, and permission can be revoked at any time.

Face Auth compared to Face ID

To users, Face Auth will feel as recognizable as Apple Face ID.

So, why not just use Face ID?

Face Auth ensures the person using World App is the same person who created the associated World ID at an orb. Face ID does not allow for this capability.

Face ID is a combination of hardware and software, so it's ultimately tied to an iPhone. With Face ID, users could have a different face associated with the device than the one they used to verify the World ID, which increases the potential for fraud. By using World ID, which is at the *app level* rather than the device level, Face Auth prevents anyone other than the person who verified the World ID from accessing it.

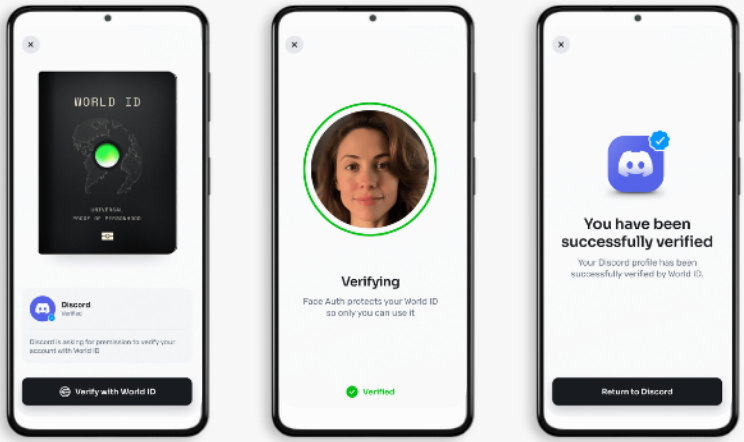

Principle 4

Transparency: Built in the open

Security experts are a skeptical lot, looking for dangers behind every line of source code. And that's a good thing, it's how they keep us safe.

It would be unfathomable to think that TFH or a few Worldcoin developers could imagine every possible fault with the protocol. Which is why Worldcoin is **built in the open**.

Audited

Worldcoin incorporates as many outside opinions and fields of expertise as possible. Cryptographers and biometrics experts are continually evaluating the source code as security auditors and consultants try to find the slightest possibility of a vulnerability. Worldcoin then publishes the results—and what it's done to address even the smallest issues.

To identify potential breaches, we also have to understand how the protocol might be used in the real world—now and in the future. That means thinking through thousands of cultural realities across the globe and incorporating these considerations into a security model that guards against misuses that may not even exist yet.

Beginning in April 2023, audit firms Nethermind and Least Authority conducted two separate security audits of the Worldcoin protocol. Specifically, the audits covered the following areas:

- Correctness of the implementation, including cryptographic constructions and primitives and appropriate use of smart contract constructs
- Common and case-specific implementation errors
- Adversarial actions and other attacks on the code
- Secure key storage and proper management of encryption and signing keys
- Exposure of any critical information during user interactions
- Resistance to DDoS (distributed denial of service) and similar attacks
- Vulnerabilities in the code leading to adversarial actions and other attacks
- Protection against malicious attacks and other methods of exploitation
- Performance problems or other potential impacts on performance
- Data privacy, data leaking and information integrity
- Inappropriate permissions, privilege escalation and excess authority

Open-Sourced and Permissionless

Making the Worldcoin code open-source and permissionless helps it fulfill three objectives. First, it exposes the network to critiques that can improve the network.

Second, it enables developers to feel confident about deploying on top of the Worldcoin protocol. It's possible that other teams might create a proof-of-personhood application on top of World ID, or find an even more usable verification method than an orb.

Finally, being open source and permissionless is essential for a decentralized project: anyone can create their own version (or "fork") of the protocol at any time, for any reason—and this is a good thing.


WORLDCOIN