



INTERNET OF PEOPLE

create your own ...

Abstract

Online identity is a mess. Despite the Internet's reputation for anonymity, it is virtually impossible to perform any action online with true privacy. At the same time, processes for willingly verifying identity or other claims are cumbersome, unintuitive, and reliant on a handful of centralized providers or imperfectly replicated versions of offline systems.

Internet of People (IOP) is building on fundamentals provided by organizations such as W3C to provide **decentralized self-sovereign identity and trust solutions** via an interdependent ecosystem of open protocols, decentralized tools, and services built with the latest cryptographic technologies and secured by the Hydra blockchain and token.

Once complete, the **IOP stack** will give users full control over their identities and their personal and professional connections. Anyone can use IOP to present **verifiable claims** without the need for a third party or compromising their privacy. This has the potential to revolutionize a wide range of industries, as well as providing access to proof of identity, credentials, and ownership to billions of currently underserved people.

Simple yet powerful tools will allow **businesses and institutions** to connect directly with their target customer base without the costs, hassle, and security and regulatory liabilities associated with harvesting and storing unnecessary data.

Among other business use cases, this will pave the way for a new kinds of **verifiable credentials** and tools for requesting and sharing them in a way which matches current user and business needs, while also opening up these markets to the huge numbers of people across the world who currently lack access to these services.

About this Document

This document serves as an introduction to Internet of People and the individual technologies and protocols created for the IOP stack. It explains the motivation and vision behind the IOP project and gives a high-level overview of the technology behind the IOP stack. You are reading **v1.1** of this document, which was created on **2019-08-06**. To access the latest version of this document visit <https://iop.global/whitepaper/>.

If you would like to learn more, visit our website <https://iop.global/>.

Contents

Abstract	2
About this Document	3
Contents	3
Introduction	4
Summary of the IOP Stack	9
Design Principles	12
Who Is the IOP Stack for?	14
How IOP Works	16
Relevant Markets	29
Development Plan	35
Getting Started	39
References	40

Introduction

Chapter Pages



Our **identity** defines who we are as individuals. Some details of our identity we keep private. Some we share with those closest to us. Others we happily share with our colleagues, our neighbours, or even the whole world. Everyone has different preferences here, and these decisions should be **ours and ours alone**.

Identity is fundamentally linked to the concept of **credentials** and **verifiable claims**. Whether it's applying for a job or a loan, buying or renting a house, travelling abroad, or simply visiting a bar, we often need a way to prove certain **claims** about ourselves, for example that we are old enough to buy alcohol or have authorization to cross a border.

Offline, the process is familiar to most of us. We have some **physical document**, such as a passport or a birth certificate, provided by some **issuer** with authority to confer the status granted by that document, such as a government or an examination board. We then show this document to an **inspector** who we need to prove a **claim** to, for example a border official when travelling to a foreign country, or a bartender when buying alcohol. This inspector then applies some process to **verify** the document, either themselves or via a third-party **verifier**. If the claim is deemed valid, **access** is then granted to the requested product, good, or service.

But online, this process is less clear. There is still no good digital analogue for physically showing a paper credential to a human verifier.

A Broken System

The challenge surrounding the provision of digital verified claims may seem unintuitive, because so many online interactions involve sharing and viewing various files and other information. But the offline claim verification process has several key features which may not be immediately apparent and which are hard to replicate online.

First, in most cases offline claim verification involves human judgment. Verifiers know what identification documents look and feel like, and they use their experience and contextual clues about the document and the interaction to decide whether it is real or not. As the likelihood of forgery and the consequences of an error increase, it is necessary to apply more checks and employ people with more experience and finely honed judgement. The doorman at a bar needs to know less about passports than a border control agent, for example. You

Introduction

Chapter Pages



need to provide more credentials and supporting documentation to apply for a mortgage than for a credit card.

Although many of these processes are now aided by technology, they still involve humans in an auditing or oversight role. These processes don't easily transfer to situations where machines are both presenting and verifying claims and the process is fully automated.

A common fix has been to try and import offline credentials into the online realm. There are various third-party verification services which lets users scan their credentials or connect to a human verifier via video link. But such systems are slow, expensive, and force users to give up control of their data.

A second major problem is secure transmission of data. In many cases of offline claim verification, the person making the claim and the person verifying it are both physically present. When documents are sent remotely, this usually occurs via a trusted transmission process, such as registered post or a company's internal logistics system. Online, this system often breaks down. It is surprisingly hard to definitively prove that the claim being transmitted and the one being received are the same, or that the people and institutions we are communicating with are who they say they are.

The current solution is to use asymmetric cryptography combined with a limited number of trusted certificate authorities. But this places an unacceptable amount of control in the hands of just a few organizations. Why should these organizations be the only ones we trust to provide and set the cost of certification services?

We need a better way to provide claim verification online.

A Decentralized Solution

Part of the problem has simply been standardization. With an agreed specification for machine-readable credentials and identification, many of the current problems could be reduced. But the bigger issue is always the reliance on a small number of trusted third parties at every stage of the process: from the issuance of credentials to their transmission and verification.

Introduction

Chapter Pages



Digital signatures based on asymmetric cryptography have powered the web and most of our digital communication for decades now. What has been missing is a way to combine this secure communication with a decentralized source of trust. Blockchain provides that source.

Some progress has already been made in this arena. The World Wide Web Consortium (W3C) specifications for decentralized identifiers (DIDs) provide a clear standard for using blockchain to secure the decentralized claim verification process [1,2,3].

But these solutions are still rudimentary, and do not provide the levels of flexibility, usability, and familiarity needed to attract a broad user base and attract people away from the centralized platforms.

To address these problems and more, the Internet of People is creating a full decentralized stack built on self-governed identity, combining a specific implementation of the W3C's DID specification on a public blockchain with a state-of-the-art consensus mechanism, as well as allowing issuance and verification of verifiable claims using user-friendly tools. Once users have established their identity, IOP will provide fully private peer-to-peer communication to connect users without the possibility of third-party interference.

This will not only streamline existing processes — it will also open up a whole host of new markets and business opportunities, and provide access to billions of people who are currently underserved or entirely neglected by existing systems.

Overlooked Billions

Despite its familiarity, the offline process isn't actually a good model to emulate. Like online, in the offline claim verification process issuing authorities act as gatekeepers for services and processes in a way that restricts choice and freedom. In addition to being vulnerable to theft, damage, and loss, identifying documents and the services surrounding them are often extremely expensive, effectively gating access to basic human freedoms such as travel and property ownership. Even more fundamentally, it isn't always clear why issuing authorities should have the control they do: for example, why should something as basic as proving your existence and chosen identity require a birth certificate issued by a hospital?

Introduction

Chapter Pages



We have become so accustomed to these systems that we no longer notice how unintuitive and restrictive they often are.

Billions of people around the world do not even have access to offline verification systems. Perhaps they cannot afford them, or their government issues or revokes them based on unfair criteria. Sometimes provision is sparse and people simply live too far away from the nearest third-party issuer or verifier.

People in these scenarios are often trapped in a vicious circle of arbitrary bureaucratic dependencies. Without certain documentation it can be impossible to enter the financial system and obtain basic services such as a bank account. But without a bank account it is impossible to have a verifiable credit history, which restricts access to loans, property, and even employment.

Billions of people are denied access to travel, financial services, and many more services which many of us take for granted, or are forced to pay excessive costs to secure them. Thanks to the rapid rise of mobile technology and the decentralized self-sovereign identity IOP provides, these people can finally access the same opportunities as the rest of the world, and businesses which operate on top of IOP will have access to a receptive and previously untapped market.

Privacy

Not all user data is alike. It sits on a spectrum from abstract to specific, depending on how easy it is to correlate that information with a particular individual. Privacy also sits on a spectrum, and most people have different attitudes to what information they are and aren't willing to share, and with whom.

Abstract, non-correlatable information includes data like being over 21 years of age, being male, liking Italian food, and so on. This information is often useful for assessing preferences, conducting broad research, and performing filtering and targeting, but cannot easily be used to identify specific individuals.

Introduction

Chapter Pages



Then there are more specific demographic variables such as date of birth, city of birth, height, university attendance, and so on. These data points are narrow enough to be very useful in all kinds of situations, but still apply to a lot of individuals, making it hard to compromise individual privacy. Note, however, that only a few data points at this level of specificity can be enough to narrow things down to a single individual [4].

Finally, there is highly correlatable information such as passport and other identification numbers, fingerprints and other biometric data, credit card numbers, and so on. This data is usually associated with just a single individual, so it is extremely hard to maintain or regain privacy if this data is compromised.

Ideally, claim verification should occur at the most abstract possible level. However, this is rarely the case. Offline, the limited number of physical credentials we have available can often force us to overshare: a passport or driving license contains far more verified information than your age, for example. And while a bartender doesn't need to know your country of birth or your home address to decide if they can legally sell you alcohol, there is no way to withhold this data when using these credentials.

Online the problem is even worse, with many interactions requiring the user to reveal far more information than is necessary. Even something as simple as logging into the WiFi in a cafe can require users to divulge their name, email address, birthdate and in some countries even passport details, even though this could easily be managed using only non-correlatable data.


Even more insidious is the "one-click" approach of logging in with an existing account, most commonly Facebook, Twitter or all-purpose apps such as LINE or WeChat. In exchange for this convenience, many users do not realize that they are sharing significant amounts of personally identifying data.

Summary of the IOP Stack

Chapter Pages



The IOP stack is a bottom-up and multi-layer architecture with different tiers and several layers of bundled projects and other technologies for different purposes and apps. Each project consists of an independent and complete core protocol which follows specific design guidelines. The goal is to write each protocol in a modular fashion and make each into a de facto standard to be used for the middle-tier programs and top-tier applications that will later run on top.



Our **Morpheus** technology allows every user to create a collection of cryptographically secure identifiers based on the W3C specification for Decentralized Identifiers (DIDs). Morpheus DIDs will be secured by the Hydra blockchain. Our blockchain has its own token, called **Hydra**, which is often abbreviated as HYD or using the symbol ₴.

Summary of the IOP Stack

Chapter Pages



Personal data never touches the blockchain in IOP, because Morpheus DIDs are fully pseudonymous. Instead, users build a **portfolio of claims** based on their Morpheus DIDs which can be selectively disclosed, forming contextually unrelated **personas**. Users can even create different personas for each of their relations, resulting in **pairwise-unique** identifiers. This collection of personas forms a coherent **digital identity**, while different personas seem completely unrelated to unauthorized third parties.

To prevent information leakage when claims are sent and verified, IOP provides fully private peer-to-peer communication via its **Mercury** protocol. All user data is encrypted by default, and Mercury uses two-way authentication based on Morpheus DIDs to ensure both parties can be sure who they are communicating with. A network of secure **Home nodes** supports this communication network and provides additional services such as secure storage.



Extending the concept of claims, users can also form relations between personas, which can be both public and private. These relations will combine to form a social network that cannot be controlled by a single entity. We call this the **Open Social Graph**.

Accessibility will be ensured via **Prometheus**, a multi-purpose app designed to be users' central hub in the world of IOP, providing them with an easy-to-use interface to access their key vault, credentials, claims, relations, and contacts.

Building on these fundamentals, any user or device will be able to use IOP technology to connect, interact, and transact freely and securely.

Summary of the IOP Stack

Chapter Pages



IOP has already made significant progress towards achieving its goals, with several key technologies in the testing phase or already complete. This also includes **Titania**, the world's first crypto-focused operating system, which is currently available for download and testing.

Once the main components of the IOP stack are complete, we plan to further develop it with technologies that will help more users access our decentralized technologies and give businesses and institutions more opportunities to provide services based on privacy, identity, and verifiable claims.

These projects include **Athena**, an educational network connecting academic institutions and learners to provide verified credentials, educational materials and scientific cooperation; **Fides**, a decentralized trust and reputation management system built on cryptographically verifiable interactions between personas; **Themis**, an arbitration and dispute-resolution system built on top of verifiable agreements and a domain-specific smart contract language which will be accessible to both lawyers and laymen; and finally **Plutus**, a governance system that allows users to shape and influence their local communities.

Design Principles

Chapter Pages



Established in 2016, Internet of People is an open-source project dedicated to ensuring everyone has access to a self-sovereign identity and the tools to transact, communicate, and express themselves online. Our technology supports a host of familiar and unprecedented use cases and business opportunities such as social networking and verifiable credentials, all without the rampant centralization, indiscriminate data commoditization, and erosion of personal privacy which are causing users to distrust and abandon today's centralized platforms.

The core components of our stack will be created using the Rust programming language, which guarantees memory safety and thread safety using compile-time checks with zero runtime overhead, preventing a whole class of the most common and devastating bugs which plague the computing industry [5,6]. However, Rust creates a high barrier of entry for developers, so we are also ensuring accessibility by providing tools to create applications on top of the IOP stack using the most familiar programming languages.

In order to promote and preserve these values, all IOP development follows a number of strict guiding principles¹:

¹ Some of the IOP design principles have been adapted from similar principles adopted by the W3C [1]

Design Principles

Chapter Pages



Decentralization	IOP technology is designed to eliminate centralized authorities, bottlenecks, and single points of failure in relation to identity and all forms of social interaction.
Self-Sovereignty	IOP technology gives users the power to directly own and control their digital identities and relations without the need to rely on external authorities.
Privacy	IOP employs privacy by design, providing tools to control users' privacy of information, including minimal, selective, and progressive disclosure of attributes or other data.
Security	IOP uses state-of-the-art technology to provide users with the security to depend on our technologies in their daily lives.
Interoperability	IOP technologies follow existing standards where possible to take full advantage of existing solutions.
Modularity & Extensibility	IOP's technologies are designed in a modular fashion to allow parallel development, extensions by third parties and to promote freedom of choice.
Portability	IOP aims to be system and network-independent, so users can use our technology on any system.
Transparency	IOP is open source and all technologies will be submitted to peer review.
Accessibility	IOP provides technologies that can be used by everyone at minimal costs.
Simplicity	IOP technology is designed to be as simple as possible to promote usability and third-party development.

Who Is the IOP Stack for?

Chapter Pages



The short answer is **everyone**.

In the long term, we hope the **IOP stack** will give everyone access to a **self-sovereign identity** that can never be revoked, as well as providing new kinds of **verifiable digital credentials** and a new relation-based rather than platform-based approach to **content delivery** and **social media**.

In addition, IOP's **user-centric** and **decentralized** approach to identity and verifiable claims can finally make identity and credentials services available to **billions of underserved and overlooked** people across the globe, while improving efficiency, security, and privacy for those ill-served by current centralized models.

These are lofty goals, but IOP also has numerous immediate and short-term applications. **Developers, companies, and investors** with the foresight to see the coming shift in power back to users can get in at the ground level of an application and platform ecosystem that, by prioritizing user privacy and security, actually provides greater efficiency and returns on invested time and resources.

While it will take some time for digital decentralized educational credentials to become widely accepted, **universities** and other **academic institutions** can already start to add blockchain-based credentials to supplement their existing approaches. This is particularly useful in developing countries, where current auditing standards are often insufficient, and indeed IOP is already working with several institutions across the world to improve the visibility and reputation of their qualifications.

Using **Morpheus** personas to create platform-agnostic relations, IOP's **open social graph** will shift the power in social media away from platform owners and back to users and content creators. The vast numbers of people from all sides of the political spectrum who have grown disillusioned with the inconsistent and unethical behavior of the major social media and content platforms can all find a home at IOP.

Who Is the IOP Stack for?

Chapter Pages



IOP's **Mercury** protocol uses two-way authentication based on Morpheus DIDs to deliver true peer-to-peer communication which cannot be disrupted or spied upon. People who live in countries with aggressive social media censorship policies will finally have a way to connect, communicate, and create freely and safely.

Titania OS is already complete and ready to try. In the medium-term we hope Titania will provide people all across the world a secure and low-cost gateway to a secure and private online world.

With these fundamentals in place, the projects at the upper levels of the stack will support the development of an ecosystem of apps and services to take IOP's decentralized identity and claim verification technologies mainstream, from universities to boardrooms and indeed living rooms around the world.

How IOP Works

Chapter Pages



IOP provides users with decentralized, self-sovereign identifiers which cannot be revoked by any third party. These unique identifiers can then be augmented with personal information by making them the subject of cryptographically verifiable claims, forming contextually separate personas that together make up the user's digital identity. On top of this secure foundation, IOP establishes a network for secure peer-to-peer communication and decentralized services.

To encourage mass adoption, it is important that establishing an identity and using the claim verification process be simple, cheap, and user friendly, all while preserving user privacy. To achieve this, IOP employs a variety of new and existing technologies.

Decentralized Identity with Morpheus¹

Morpheus uses decentralized identifiers (DIDs) based on the W3C standard. DIDs are a form of URL and resolve to DID Documents. These DID documents contain information about the DID, such as a list of public keys or other information which can be used to verify the identity of someone trying to act on behalf of this DID. DID documents may also contain information on any service endpoints the controller of the DID wants to advertise. Crucially, DID documents contain no personal data and DIDs are therefore pseudonymous identifiers. Users may register any number of unique decentralized identifiers (DIDs) on the Morpheus DID registry.



¹ Terminology and diagrams in the following two sections are adapted from W3C documents on decentralized Identifiers and verifiable claims [4,5,6]

How IOP Works

Chapter Pages



The Hydra Blockchain as Decentralized Root of Trust

If we want to provide verifiable identity without relying on a central trusted authority, we need a decentralized root of trust. There are several options here, but blockchains and other decentralized ledger technologies that have a consensus mechanism including economic incentivization for maintainers have a specific advantage over other decentralized forms of storage: Financial gain is a strong incentive for node owners to adequately maintain agreement on global state, protect the stored data against accidental or deliberate corruption, and to accept any and all transactions users might make.

The Morpheus DID registry will be supported by the Hydra blockchain. A DID is registered using a transaction very similar to a standard cryptocurrency transaction, including a fee to register and store the DID document.

Because DIDs are pseudonymous identifiers, personal data never comes near the chain. Instead, all personal information is encrypted by default and stored on the user's end device or backed up on another storage solution of the user's choice. If users want to publicly disclose information about a DID, they may do so by including a service endpoint in the DID document.

Using this approach provides multiple benefits. First, it prevents clogging the chain with repeated and unnecessary changes to DID documents. Second, history has shown that every encryption protocol is eventually broken. For this reason it would be irresponsible to store personal data in an immutable and eternal database.

As a result, DID documents will only change rarely for regular users — generally when DID documents are first created or when keys are lost or compromised. This keeps barriers to entry low for users. However, the transaction flow for companies and other institutions will be very different, as they will likely need to manage hundreds or thousands of service endpoints and keys for different DIDs. This steady flow of transactions will keep the Hydra token liquid and provide reliable monetary incentives for node maintainers.

How IOP Works


Chapter Pages



Identity Verification

The process of identity verification in the offline world can be described in an abstract way:

- An identity holder requests a benefit from an inspector.
- The inspector asks for proof of identity.
- The holder of the identity shows their identity credential to the inspector.
- The inspector sends this credential to a verifier for verification.
- The verifier checks the credential against a registry which stores pre-registered credentials.
- The verifier returns the result to the inspector who either approves or denies the holder.



These roles can be held by people or machines. In the case of airline travel using a biometric passport, for example: the **inspector** is the border agent, the **verifier** is the passport scanning machine and the **registry** is the centralized database of passport records. Multiple roles may be fulfilled by the same actor. When using the same passport to buy alcohol at a bar, for example, the inspector, verifier, and registry roles are all fulfilled by the bartender (with the registry being a mental checklist of what a valid passport looks and feels like).

How IOP Works

Chapter Pages



Let us now try to translate this example to the digital world. To keep things simple, we will assume that the inspector is also performing verification. (When the two roles are separate, there is simply an extra step where the inspector sends a request to the verifier and waits for a response.)

To prove they control a DID, a user sends the DID to an inspector/verifier, along with a cryptographic signature created using a private key. The inspector/verifier then uses the DID to resolve the DID document in the Morpheus registry. The DID document contains all public keys associated with this DID, so the verifier can check the cryptographic signature and verify that the user indeed controls this DID.

Verifiable Credentials without Central Authority

Proving control over a randomly assigned DID is of very limited use without a way to associate it with real-world information. We would like to be able to make **verifiable claims** about ourselves, both on- and offline. By creating various DIDs, users can then create multiple, fully separate sets of individual claims which anyone can verify without the need for a centralized trust source.

It is perhaps most intuitive to think about these claims in terms of familiar data points (e.g., name, date of birth, graduated with a Master's in Computer Science from Stanford, etc.) drawn from familiar offline credentials (e.g., birth certificate, passport, university degree), but the range of possible claims is far, far broader than this.


Unfortunately, in the virtual world it is almost impossible for users to prove any piece of personal information on their own, apart from cryptographic proofs of key control. When making a personal claim to an **inspector/verifier**, it is necessary to appeal to a trusted third party.

How IOP Works

Chapter Pages



For digital verifiable claims, the trusted source comes in the form of an **issuer** who issues claims about the user in the form of **credentials**. Users can apply to issuers to issue claims about them which they can later present to an inspector/verifier. Once a credential has been issued, the holder can use the claims it contains to request access to goods and services. On receiving a claim, the appropriate inspector/verifier checks that the claim is genuine. If it is, they then grant the holder access to the agreed benefit.




As an intuitive example of this trust model, imagine you are applying for a job and need to demonstrate that you have the necessary university degree. It is impossible to demonstrate this on your own, but your prospective employer will presumably trust the authority of the university.

How IOP Works

Chapter Pages



Note that this trust model is completely different from being reliant on a centralized root of trust. There, everyone is forced to choose from a single or very limited number of trusted sources. In a decentralized trust model, parties are always free to choose any source of trust which they can agree upon.




As mentioned before, it is vitally important that personally identifying information never comes anywhere near the blockchain; therefore, the information that users want to make claims about must be stored separately and handled differently from the DIDs and DID documents which make up the Morpheus DID Registry.

How IOP Works

Chapter Pages



Nevertheless, as shown in the figure below, the DID Registry provides the necessary root of trust, as it is the gateway to the data needed to verify DID ownership and credentials and claims based upon them.



How IOP Works

Chapter Pages



The figure below shows an example claim which can be used to prove that the user holds a Bachelor's Degree in Mechanical Engineering. One significant advantage of this system is that it can be used to minimize information exchange. In the example below, the issuer probably knows the holder's graduation date and has more information like the topic of the user's thesis. However, the inspector/verifier is only interested in whether the holder has a degree or not. The issuer is able to answer this request without disclosing additional information about the credential holder.



Personas

The claim verification system described above is extremely useful, but it can also quickly become hard to manage as the number of claims and credentials increases. In IOP, credentials are managed according to the intuitive concept of **personas**. A persona is the collection of claims about a single DID. These claims can safely be disclosed together,

How IOP Works

Chapter Pages



without contaminating other DIDs. For example, users might choose to have one persona for interacting with friends and family, another for professional contacts, and another for online dating. Personas can be cryptographically linked and switched between as the user chooses, without their underlying connection being discoverable by a third party. The claims that make up a persona can of course still be disclosed individually, providing users with fine-grained control about which parts of their personal information they want to share.

Secure Peer-to-Peer Communication and Storage

In a privacy-focused system such as IOP, it is not enough to provide secure identity and claim verification. It is also crucial to ensure that no data is leaked when users issue, present, or verify claims or communicate in any other way.

In addition, while storing user data locally and encrypted by default promotes privacy, this is insufficient in many cases. First, local storage space is limited, especially as many users will be interacting with IOP via smartphones. Second, end devices are not always online, regularly change networks, and are generally unreliable with limited storage options. Third, many users will want to access their collection of claims from multiple devices.

Additionally, the acts of registering and updating DID documents and issuing and verifying claims rely on access to a complete and sufficiently current version of the Morpheus DID registry. This might not be feasible for some devices, which can instead rely on the network of Home nodes that will store the entire history. We foresee a new generation of service providers that maintain the full database as a service, similar to current web wallets, explorers, et cetera.

To accommodate these usability and accessibility requirements, IOP's Mercury protocol provides encrypted, two-way authenticated peer-to-peer communication to maximize privacy and security. Generating truly private peer-to-peer connections between users is a particularly difficult problem, particularly if one or both users is on a smartphone. Modern Internet architecture locks users behind multiple layers of Network Address Translation (NAT) which makes it impossible to address users directly, especially when trying to make contact for the first time.

How IOP Works

Chapter Pages




Mercury uses a network of so-called Home nodes to punch through these layers and open up a direct channel of communication. When two users want to communicate, the first user begins by directly contacting their trusted Home node(s). They then communicate with other Home nodes until they find one which knows how to contact the second user. This multi-hop route is then collapsed to provide a direct channel of communication.

In this way, Mercury acts similar to a decentralized and highly secure cellular network, where Home nodes take the roll of cell towers and devices (mobile or stationary) are phones. In contrast with mobile networks, nodes are not selected automatically based on location and signal strength, but picked manually based on trust and reliability.

Mercury is a decentralized network because all Home nodes have the same role: a node can join or leave the network any time. The same applies for clients. This makes the network extremely resilient as there is no main point of attack and the network dies only with the last node.

In addition to maintaining the communication network and supporting the Hydra blockchain, these Home nodes will be able to provide secure data storage and claim verification services to users.



How IOP Works

Chapter Pages



Advanced users can operate their own Home node. Home nodes require minimal resource, so can be run on most devices with low CPU/memory and reasonable network requirements. **Titania OS** will be the perfect starting point to set up your own Home node as it comes preconfigured with the full IOP stack. After setting up Titania with just a few clicks, it can act as a repository for user claims, personal storage space and as a gateway to the Mercury network.

The functionality of Mercury extends far beyond the IOP stack. Combining Mercury's revolutionary concept of establishing connections based on an identifier:protocol scheme (instead of the hostname:port scheme that today's applications use) with standardized profile and message formats allows app developers to create decentralized applications (dApps) that are not reliant on a centralized database of users.

The average user will simply install applications built on IOP's decentralized technology to a device such as a smartphone or PC or open a web-based dApp. These applications will rely on Prometheus as a dependency, which will bring a secure vault for keys and claims, background services, and a convenient UI to manage personas, claims and relations. After initial setup, users choose which persona to use with each dApp and choose one or more Home Nodes to connect to. Mercury's unique architecture then allows connections beyond the confines of a dApp, relieving developers of the need to laboriously create a critical user base and allowing users to choose which dApps best fit their needs.

Broader IOP Functionality

While the ability to verify identity and claims without relying on a centralized root of trust is obviously useful, the preceding sections' focus on the technical minutiae of individual interactions can make it hard to appreciate the power of these technologies when scaled up to entire organizations, communities, or whole societies.

The upper layers of the IOP stack contain projects designed to meet broad societal needs in a decentralized fashion. While these all have decentralized identity and claim verification at their core, they expand on these basic concepts to provide powerful decentralized services such as contract resolution and arbitration, reputation, governance and more.

How IOP Works

Chapter Pages



Athena

Athena applies IOP's tools and technologies to education. The most obvious application is the provision of a new kind of trusted and verifiable educational credential, but academic institutions face a whole host of problems which can be addressed through secure decentralized identities and verifiable claims, from attendance to content and syllabus provision.

Notably, institutions which provide learning at a distance and online materials can easily manage course content and access, and Morpheus DIDs can be used to secure online testing environments. In this way, IOP can improve access and educational opportunities for millions of currently underserved people.

Athena will also facilitate research and collaboration between institutions. Verifiable claims can be used to improve processes for data collection, sharing, and analysis, and the terms of any collaboration can be easily managed via on-chain smart contracts.

Fides

A wealth of verifiable claims can lead to information overload. Users will need a way to quickly assess whether they want to communicate or transact with each other.

Fides will provide users with tools to navigate a system based around verifiable claims. As more public and private claims become available about a persona, making decisions about the trustworthiness of that persona will become more complex. This is especially true if that persona is linked to a business or major service provider with thousands of customers and associated claims. One major element of this is decentralized trust and reputation tools based on cryptographically verifiable interactions between personas. For familiar reputation and ranking systems such as Uber ratings or TripAdvisor reviews, this has the obvious benefit of linking reviews and feedback to verifiable identity, reducing problems like spam and sock-puppeting which plague existing reputation systems. But Fides also drastically broadens the kinds of reputation data which can be collected and used, making reputation systems much more powerful and reliable.

How IOP Works

Chapter Pages



For example, a user leaving a review in an online marketplace can prove that they genuinely bought the product being reviewed, without giving up any other personally identifying data. Reputation data itself can also be used as a basis for verifiable claims, for example in loyalty or VIP programs.

Beyond this, Fides will provide tools to manage relations between trusted personas. In addition to broad social media applications such as the **Open Social Graph** and content platforms, this includes more specific tools such as providing ways to organize clubs, events, or even employee rosters.

Themis

Themis provides tools to assist with arbitration and dispute management, including an intuitive language for smart contracts which is accessible for both lawyers and laymen.

One major problem with existing smart contracts is the difficulty of accessing and verifying the off-chain data which the contract is contingent on. The current solution is to use oracles, but they are clumsy and cumbersome. Verifiable claims provide a better way for on-chain contracts to interact with off-chain data.

Disputes will still arise which can only be resolved through existing systems such as courts, but Themis will help users design smart contracts to be as robust as possible and maximize the chance of reaching a resolution between the contracting parties.

Plutus

Finally, **Plutus** is a project dedicated to safely and prudently scaling up IOP's technologies to address larger-scale needs, for example in large corporations or local and city governance. Verifiable claims and decentralized identity can revolutionize our approach to complex and contentious problems such as managing eligibility for healthcare or insurance schemes, or running secure and auditable votes without gating access or exposing voter data.

The following section takes a closer look at the various markets where IOP technologies can be leveraged for improvement and disruption.

Relevant Markets

Chapter Pages



The data and identity industries have become huge business, touching every aspect of our modern digital lives.

But the cracks are beginning to show.

Some of this can be attributed to the business practices of the centralized giants such as Facebook, Amazon, Uber, and others, but many of the problems run much deeper: they are inherent in every centralized approach to identity and personal data, and cannot be fixed without a totally new model.

IOP provides that new model, one which benefits both users and businesses. Users get more control, stronger privacy, and can be directly rewarded for sharing their personal data. Businesses get easier access to a greatly expanded potential user base and also receive a bigger slice of the potential pie, instead of subsisting on the crumbs given out by the centralized giants.

The quality of data and datasets will also increase: instead of indiscriminately harvesting data and then curating it, with IOP data is requested on an as-needed basis, making datasets more focused and easier to manage, analyze, and use. Individual data is also easier to assess thanks to claim verification. Finally, by switching to a model where users willingly and actively share data rather than having it harvested in the background, users will have more incentive to provide accurate and helpful responses to data requests.

Here are just some of the markets which IOP can disrupt.

Social Networking

It's taken barely twenty years for social networking to rise from nothing to encompass the whole world. Facebook currently has more than two billion monthly active users worldwide, and had a total market capitalization of over half a trillion dollars at its peak in early 2018 [\[10\]](#).

Relevant Markets

Chapter Pages



But a recent string of scandals have highlighted the significant flaws in Facebook's "go fast and break things" approach. Regulators are circling, market valuation is dropping, and user engagement is beginning to fall rapidly. Businesses are also waking up to the problems with Facebook's model: the recent reports on Facebook's notoriously opaque analytics tools showed that Facebook is not nearly as useful for engaging and retaining customers as they claimed.

IOP and the Open Social Graph inverts the social network business model by giving control back to the users. In doing so, it creates a fully transparent environment for businesses to directly engage with their customers.

Digital Advertising

The digital advertising market is huge. In the US, digital advertising spends are predicted to eclipse all other forms of advertising by 2021. But as with so many online industries powered by data, people are seriously questioning whether the current models are benefiting anyone but the centralized platforms.

Most people want to be able to find products and services that will interest them, and are happy to learn about these directly from companies. But the underhand way in which identifying data is gathered and used for targeting fosters suspicion and resentment. Businesses feel compelled to engage in digital advertising, but opaque analytics tools make it impossible to tell whether an advertising spend is cost effective.

Centralized platforms have a virtual monopoly in the advertising space, giving them free reign to set uncompetitive rates. Even though the industry is powered by user data, users receive very little in return. Users are increasingly dissatisfied with the model, which may explain why click through rates are less than 0.1% and almost half of all paid online ads are never even seen by anyone [11].

IOP removes the need for a central platform authority like Facebook, allowing users and companies to connect directly. Because data on IOP is only ever willingly shared, targeted advertising becomes simpler, cheaper, and much more effective.

Relevant Markets

Chapter Pages



Cybersecurity

Harvesting personal data isn't just bad for users: it can also become a massive liability for companies, who may not have the expertise to properly store and secure that data. Data breaches are a daily occurrence, despite yearly cybersecurity spending approaching \$100bn (as estimated by Gartner [12]).

IOP provides significant cybersecurity benefits over existing centralized approaches. At a basic level, making local storage and full user control the default for all data significantly reduces the risk of hacking or accidental exposure of data.

Beyond that, IOP provides numerous opportunities for companies wishing to offer security as a service. For example, companies could provide secure backup for users' claims and credentials or give them the ability to share and access credentials from the cloud. Unlike current cloud-based approaches, data security would not be compromised as all data would be encrypted before transmission to the service provider.

Identity and Access Management

Many companies and other organizations need to manage identity and access for employees and visitors to their sites, with yearly global spending in excess of \$8bn [13]. But current approaches are extremely expensive, inflexible, and often exhibit glaring security vulnerabilities.

Using IOP's identity protocols, companies would be able to issue identity cards and adjust access rights on the fly, as well as easily granting and revoking visitor access thanks to standardization.

This concept can easily be extended from physical to digital access, for example managing access rights to shared company drives and other digital resources from users' own devices.

Relevant Markets

Chapter Pages



Data Integration

Despite the recent rush to harvest as much data as possible, as much as 90% of it languishes in data silos, unused, unanalyzed, and even unverified. Finding ways to share data between these siloes has become big business, with an estimated yearly revenue of \$12.24 billion [14].

In addition to improving privacy, IOP actually makes it far easier to gather, collate, and analyze the data which users are willing to share. Thanks to IOP's standardized approach to data storage, companies offering data integration services will find it simple to convert and merge existing databases.

Peer-to-Peer Content Delivery

In modern media, digital content is king. More than 66% of users report that they get at least some of their news from platforms such as YouTube, Twitch, and Medium, and digital content creators can earn millions of dollars a year [15].

But existing models are strained to breaking point. YouTube and Twitter are unable to consistently apply their own rules, and their recommendation and monetization algorithms are subject to massive abuse. Social media lifecycles are also incredibly short, especially for apps. Content creators spend years building up followers across various platforms, only to have all that hard work erased when an app falls out of fashion or a platform's terms and conditions suddenly and arbitrarily change.

With IOP, control of content provision and delivery will be returned to users and creators. Creators and their followers will be directly linked via Morpheus personas, and this relationship can persist across any number of separate platforms. Instead of relying on centralized platforms for content curation, users can easily and directly decide whose content they want to see and whose to filter out. The resulting systems will be freer and more democratic, and the creator/follower relationship will be preserved even as social media trends come and go.

Content can be shared directly using the Mercury protocol, or via apps and sites built on top of the IOP stack. These platforms will be secure and censorship resistant. They will also provide freer and fairer monetization options, as there will be no central authority holding the purse strings.

Relevant Markets

Chapter Pages



Credentials

Online learning is poised to revolutionize the way we learn and open up educational opportunities to billions of currently underserved people around the world. But qualifications earned via e-learning platforms currently have a poor reputation, even when affiliated with established brick and mortar institutions. Part of this is resistance to change, but it's also undeniable that online credentials are hard to track and easy to forge for both unscrupulous users and institutions alike.

IOP's technology stack will allow universities and students to track and prove attendance, grades and degrees as well as less tangible factors such as quality of education and student satisfaction using mutually provided verifiable claims. Future employers can easily verify grades and credentials with the issuing entity without the danger of personal user data falling into the wrong hands or companies overreaching their boundaries and asking for more data than they really need.

Reputation

The rise of big data and advanced analytics has allowed reputation systems to expand from traditional industries such as insurance and credit ranking to become a multibillion dollar industry that serves multiple sectors. From ranking systems such as Uber's star rating system to more generalized feedback mechanisms, reputation systems are an invaluable tool for matching users as well as providing ways to identify and filter out malicious or unwanted behaviours and actors, particularly in systems designed for zero or minimal centralized oversight.

But current practices are beset with problems. Today's centralized reputation systems rely on enormous datasets, where individual pieces of identifying data are collated to form a complete picture of all of a user's activities. This is a treasure trove for hackers, and indeed some of the largest data breaches in recent years have been attacks on companies providing or using reputation systems.

Relevant Markets

Chapter Pages



Even more ominously, governments are turning to reputation systems to rank their citizens and control their access to goods and services and even opportunities such as jobs and education. In addition to being a massive invasion of privacy and liberty, these systems are often secret, arbitrary, and extremely poorly secured and designed [16].

IOP provides the benefits of reputational systems while eliminating the most severe security and privacy risks for both users and service providers. The **Fides** project will allow companies which need rating or feedback mechanisms — or who wish to provide these to other companies — to do so with ease.

IOP's verifiable digital claims also offer numerous avenues for improving rating and reputation systems. Current systems rely almost solely on user opinion, often employing 5-star systems which quickly degrade into binary rating systems where anything less than 5 stars is considered a failure. Combining user input with verifiable claims about good or bad behaviours will make reputation systems more flexible, reliable, and ultimately valuable.

Regulatory Technology

For many businesses, complying with ever-changing data regulations is costly and time consuming, especially when operating across multiple jurisdictions. Companies must find a way to gather, process, and then securely store user data such that they can prove compliance — often for audits that occur many years after the fact. For many smaller businesses, these costs are so high that they are forced out of the market.

With IOP's approach to identity and verifiable claims, regulatory compliance is streamlined and risk is reduced for both users and companies. Because data tends to stay under full user control, it is easy to prove that it has been handled properly.

At the same time, by completely separating personal data from the ownership of a digital identity, IOP support far more secure and stringent data protection rules than are currently implemented in most countries around the world.

Development Plan


Chapter Pages



Internet of People has been under development in one form or another for more than two years now, and some of our technologies are already complete and ready to use. Others are entering the final development and testing stages, some are being created right now, and some projects are still in the design stage.


The different technologies in the IOP stack gain their value through a symbiotic relationship, building on top of each other and interlocking almost seamlessly. Therefore, we are not creating projects in strictly linear order. The roadmap on the next page describes the order we expect the most relevant feature sets of each project to be available to users and developers.

The current development status, including regular updates directly from our development team, can be found at <https://iop.global/>.



Development Plan

Chapter Pages



Titania

The core operating system is complete and ready to use. Titania provides a fully dockerized environment for secure deployment of apps and services. Setup and management is easy thanks to an intuitive web interface. The next step is to package the software people want and need for Titania's secure environment, including staking wallets for the most popular PoS coins. A Titania dApp store is under development, and the complete IOP stack will be integrated by default. Once the Mercury network is functional, users will be able to create their own Home nodes and provide services to their peers.



Mercury

A prototype of the Mercury communication protocol has been published on our GitHub repository, with the Home and communication protocols close to completion. A few key features still need to be implemented before Mercury is ready for production use, but the underlying concepts have been clearly demonstrated. Mercury provides direct p2p communication and a decentralized network for peer discovery that not only allows users to communicate securely, but also enables developers to create powerful dApps aimed at connecting users. Once Morpheus enables ID management, Mercury will be extended with user personas and networking features to create the Open Social Graph.



Prometheus

The prototype of our Prometheus application can be downloaded from our website. Prometheus already allows access to a basic version of DID and claim management that will enable the Open Social Graph and includes an abstract key vault that can later be used by any application in need of public-key based cryptographic functionality. In the future, Prometheus will be extended into an extensive and intuitive graphical user interface to provide identity and privacy management, and will provide a comprehensive SDK so developers can use IOP technologies, including Mercury communications and Morpheus DIDs and claims, in the applications they create.

Development Plan

Chapter Pages



Morpheus

A minimal example of Morpheus' self-sovereign identity management system is included in the Prometheus Prototype. The first application of verifiable claims will be user relations that can be extended with metadata. In the future, the Morpheus DID registry will be secured by the Hydra blockchain and pairwise-unique identifiers, and fine-grained claim templates will allow users to manage their digital identity in an intuitive way, aiming for maximum privacy and minimal disclosure. The verifiable claims created using Morpheus will create a global web of trust.



Athena

We are currently defining the basics of educational credentials built on Morpheus' verifiable claims. These will give employers, students, graduates, and universities a standardized way to issue and verify attendance, credentials, and more. Because of their unique relation to thousands of people and businesses alike, universities are prime candidates for widely accepted trust anchors of decentralized identity management. Combining the features of Morpheus with Mercury's unique communication protocol will allow researchers and teachers to manage access to resources and collaborate securely on a global scale.



Hydra

Hydra provides a completely open and permissionless consensus mechanism, allowing all users to participate and benefit by staking tokens. Hydra users use their stake to elect 53 delegates from their midst to forge new blocks and create new tokens. These delegates share a large percentage of the block rewards with their voters. Delegates also earn voter share by contributing to the project or by helping new users. Everyone can become a delegate and earn rewards if they gain the support of the community. To forge new blocks, delegates need only to be online and run a reliable node. Voters can earn rewards while keeping their stake offline and secure. Hydra provides final confirmation after each block, which makes Blockchain rollbacks impossible, thus protecting the Morpheus DID register from manipulation. In the future, the functionality of the blockchain will be tightly integrated into the Prometheus key vault and optimized for the purpose of DID management.

Development Plan

Chapter Pages



Fides

Fides is our version of a bottom-up reputation system, decentralizing trust, linking feedback and reviews to verifiable identities, and giving people the tools to navigate IOP's ever-growing collection of verifiable claims. Services that aggregate multiple claims into comprehensive certificates of trustworthiness are just one of the many possibilities Fides will provide. Fides will also enhance Athena by allowing students to rate their universities' quality of education, the comprehensiveness of individual courses, and much more. Verifiable claims will be extended further to include receipts, loyalty programs, and much more, and Mercury will be extended to allow for organization of social events.



Themis

Themis will provide a domain-specific language aimed at laymen and lawyers. By feeding off-chain data into the system using verifiable claims, Themis imports concepts such as documents signed in the presence of witnesses into the digital world. This will open the door for digital employment agreements and other legal contracts that are both easy to understand and unambiguous. Themis aims to close the divide between code and legal language, creating smart contracts that can be both run by machines for automatic arbitration as well as serve as legal documents in a court of law in the case of unresolvable disputes.



Plutus

The final step will be to prudently and safely scale the technologies created by the Internet of People into a system for local governance. Verifiable claims and decentralized identity can revolutionize our approach to complex and contentious problems such as managing eligibility for healthcare or insurance schemes, or running secure and auditable votes without gating access or exposing voter data. Plutus will allow users to be heard and shape their communities according to their needs.

Getting Started

Whether you're an investor, a developer, or just someone who's interested in the possibilities of decentralized technology, we'd love to hear from you.

Visit us at <https://iop.global/>, where you can learn more about Internet of People and how to get involved.

You can download Titania OS at <https://github.com/Internet-of-People/titania-os/releases>.

You can download a pre-release version of Prometheus at <https://github.com/Internet-of-People/prometheus-electron/releases>.



<https://iop.global/wallets/>



<https://iop.global/exchanges/>



<https://discordapp.com/invite/gsRKp6T>



<https://t.me/loPofficial>



<https://github.com/Internet-of-People>



https://twitter.com/loP_community



<https://www.youtube.com/channel/UCmvh91s11TTwlJuo8-y7frQ>



https://www.reddit.com/r/loP_Community/

References

1. W3C Credentials Community Group, Decentralized Identifiers, URL: <https://w3c-ccg.github.io/did-spec/>
2. W3C Credentials Community Group, Proposed Verifiable Claims Architecture, URL: <http://w3c.github.io/webpayments-ig/VCTF/architecture/>
3. W3C Credentials Community Group, Verifiable Credentials Data Model, URL: <https://www.w3.org/TR/verifiable-claims-data-model/>
4. L. Sweeney, Simple Demographics Often Identify People Uniquely. Carnegie Mellon University, Data Privacy Working Paper 3. Pittsburgh 2000, URL: <https://dataprivacylab.org/projects/identifiability/paper1.pdf>
5. The Rust Programming Language, URL: <https://www.rust-lang.org>
6. Alex Gaynor, The Internet Has a Huge C/C++ Problem and Developers Don't Want to Deal With It, URL: https://motherboard.vice.com/en_us/article/a3mgxb/the-internet-has-a-huge-cc-problem-and-developers-dont-want-to-deal-with-it
7. Yossi Gilad, Rotem Hemo, Silvio Micali, Georgios Vlachos, Nickolai Zeldovich, Algorand: Scaling Byzantine Agreements for Cryptocurrencies, URL: <https://dl.acm.org/citation.cfm?id=3132757>
8. Jing Chen and Sergey Gorbunov and Silvio Micali and Georgios Vlachos, ALGORAND AGREEMENT: Super Fast and Partition Resilient Byzantine Agreement, URL: <https://eprint.iacr.org/2018/377>
9. Jing Chen, Silvio Micali, Algorand, URL: <https://arxiv.org/abs/1607.01341>
10. Jay R. Corrigan, Saleem Alhabash, Matthew Rousu, Sean B. Cash, How much is social media worth? Estimating the value of Facebook by paying users to stop using it, URL: <https://doi.org/10.1371/journal.pone.0207101>
11. App Nexus, The Digital Advertising Stats You Need for 2018, URL: https://www.appnexus.com/sites/default/files/whitepapers/guide-2018stats_2.pdf
12. Tech Crunch, Global Cybersecurity Spending to Grow to \$86.4bn in 2017, says Gartner, URL: <https://techcrunch.com/2017/08/16/global-cybersecurity-sending-to-grow-7-to-86-4bn-in-2017-says-gartner/>

References

13. Orbis Research, Identity and Access Management-Global Market Outlook (2016-2022),
URL: <https://www.orbisresearch.com/reports/index/identity-and-access-management-global-market-outlook-2016-2022>
14. Markets and Markets, Data Integration Market by Component (Tools and Services), Business Application (Marketing, Sales, Operations, Finance, and HR), Deployment Model, Organization Size, Vertical, and Region - Global Forecast to 2022,
URL: <https://www.marketsandmarkets.com/Market-Reports/data-integration-market-61793560.html>
15. Markets and Markets, Commercial P2P CDN Market by Content Type (Video and Non-video), Solution (Web Performance Optimization, Media Delivery, and Cloud Security), Service, End-User Segment (Consumer and Business), Vertical, and Region – Global Forecast to 2023,
URL: <https://www.marketsandmarkets.com/Market-Reports/commercial-p2p-cdn-market-187880203.html>
16. Devin Coldewey, When Surveillance Meets Incompetence,
URL: <https://techcrunch.com/2019/02/19/when-surveillance-meets-incompetence/>



<https://iop.global/>