Microsoft

Decentralized Identity

Own and control your identity

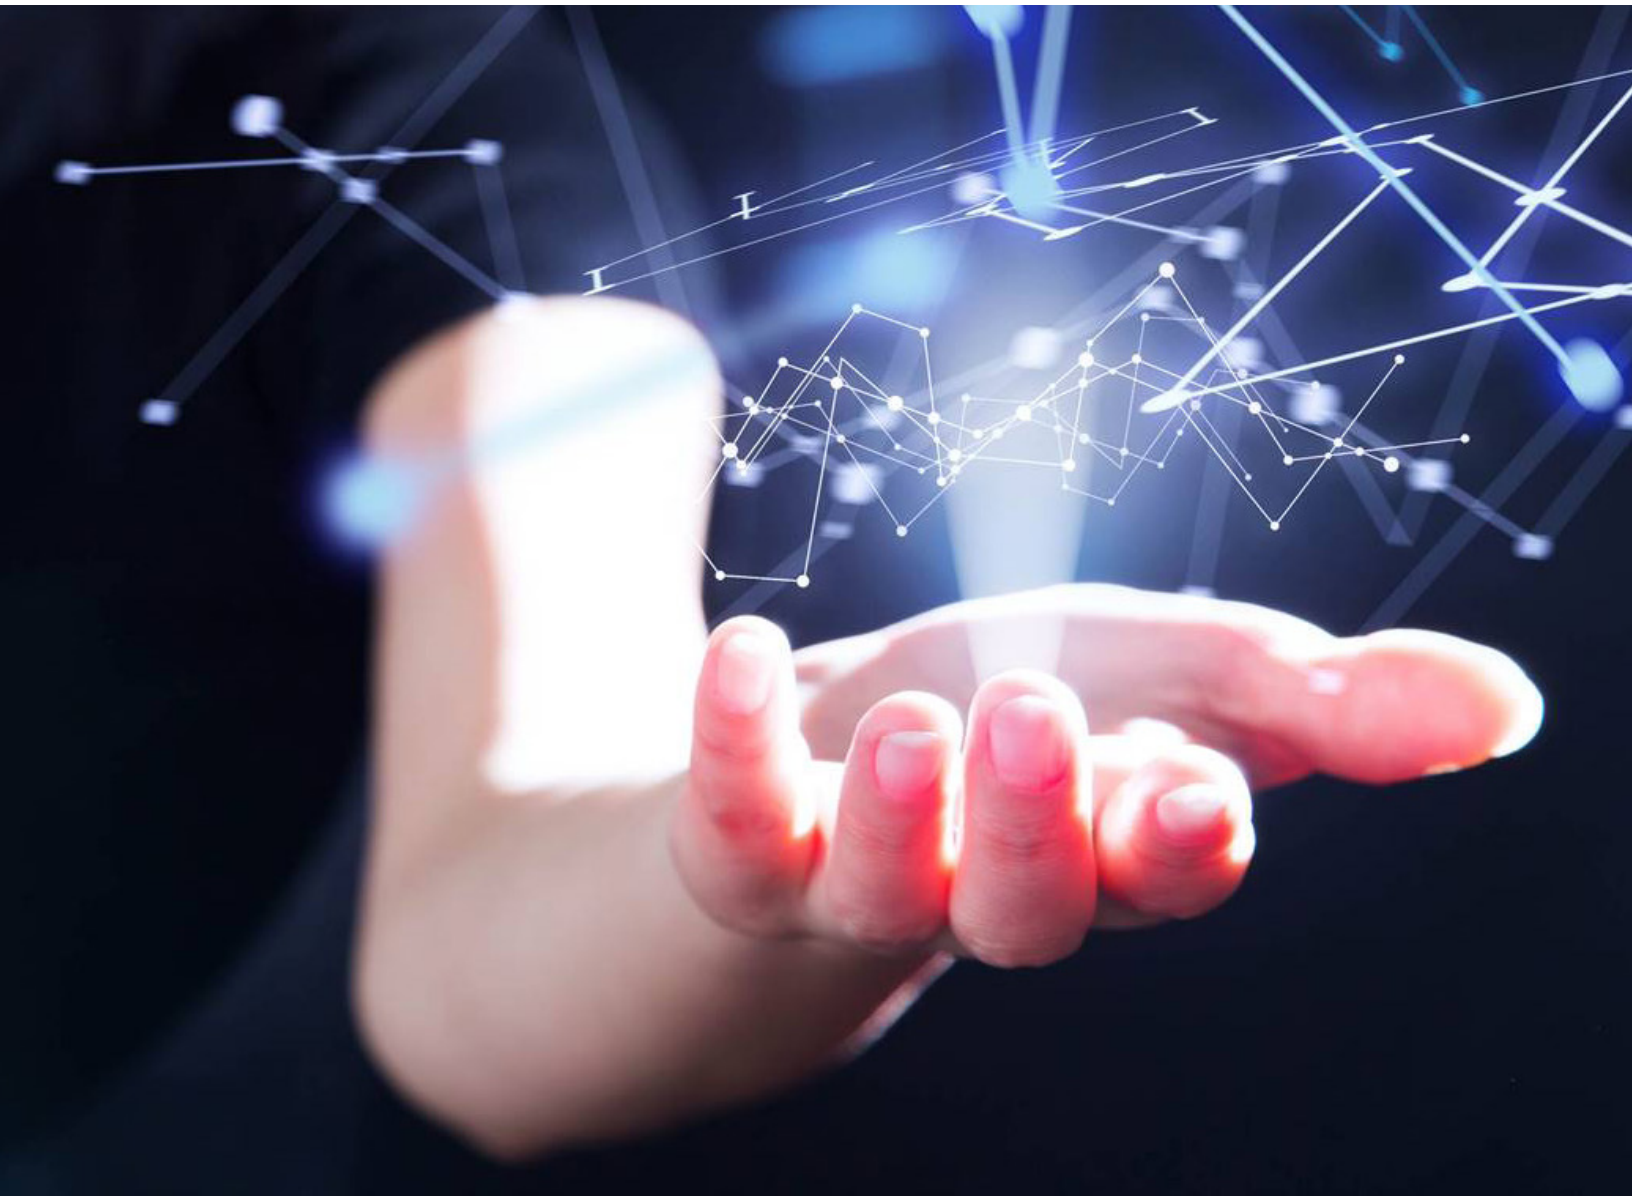

Contents

Executive summary

Our digital and physical lives are increasingly linked to the apps, services, and devices we use to access a rich set of experiences. This digital transformation allows us to interact with hundreds of companies and thousands of other users in ways that were previously unimaginable.

But identity data has too often been exposed in breaches, affecting our social, professional, and financial lives. Microsoft believes that there's a better way. Every person has a right to an identity that they own and control, one that securely stores elements of their digital identity and preserves privacy. This whitepaper explains how we are joining hands with a diverse community to build an open, trustworthy, interoperable, and standards-based Decentralized Identity (DID) solution for individuals and organizations.

Each of us needs a digital identity we own, one which securely and privately stores all elements of our digital identity.

This self-owned identity must seamlessly integrate into our lives and give us complete control over how our identity data is accessed and used.

Why we need Decentralized Identity

Today we use our digital identity at work, at home, and across every app, service, and device we engage with. It's made up of everything we say, do, and experience in our lives—purchasing tickets for an event, checking into a hotel, or even ordering lunch. Currently, our identity and all our digital interactions are owned and controlled by other parties, some of whom we aren't even aware of.

The status quo for users is to grant consent to numerous apps and devices, which warrants a high degree of vigilance of tracking who has access to what information. On the enterprise front, collaboration with consumers and partners requires high-touch orchestration to securely exchange data in a way that maintains privacy and security for all involved.

We believe a standards-based Decentralized Identity system can unlock a new set of experiences that empowers users and organizations to have greater control over their data—and deliver a higher degree of trust and security for apps, devices, and service providers.

Microsoft's strategy for Decentralized Identity

Microsoft's mission is to empower every person on the planet to achieve more.

Microsoft cloud identity systems already empower developers, organizations, and billions of people to work, play, and achieve more, but there's so much more we can do to create a world where each of us, even in displaced populations, can pursue our life goals, including educating our children, improving our quality of life, and starting a business.

To achieve this vision, we need to augment existing cloud identity systems with one that individuals, organizations, and devices can own so they can control their digital identity and data. This self-owned identity must seamlessly integrate into our daily lives, providing complete control over what we share and with whom we share it, and—when necessary—provide the ability to take it back. Instead of granting broad consent to countless apps and services and spreading their identity data across numerous providers, individuals need a secure, encrypted digital hub where they can store their identity data and easily control access to it.

Lead with open standards

We're committed to working closely with customers, partners, and the community to unlock the next generation of Decentralized Identity–based experiences, and we're excited to partner with the individuals and organizations that are making incredible contributions in this space. If the DID ecosystem is to grow, standards, technical components, and code deliverables must be open source and accessible to all.

Microsoft is actively collaborating with members of the Decentralized Identity Foundation (DIF), the W3C Credentials Community Group, and the wider identity community. We're working with these groups to identify and develop critical standards. We're developing an open source DID implementation that runs atop existing public chains as a public Layer 2 network designed for world-scale use. The purpose of this implementation is to establish a unified, interoperable ecosystem that developers and businesses can rely on to build a new wave of products, applications, and services that put users in control.

How does Decentralized Identity work?

Today, the digital representation of a user's identity is a mix of data fragmented across many apps and services.

A new form of identity is needed, one that weaves together technologies and standards to deliver key identity attributes—such as self-ownership and censorship resistance—that are difficult to achieve with existing systems.

To deliver on these promises, we need a technical foundation made up of seven key innovations—most notably, identifiers that are owned by the user, a user agent to manage keys associated with such identifiers, and encrypted, user-controlled datastores.

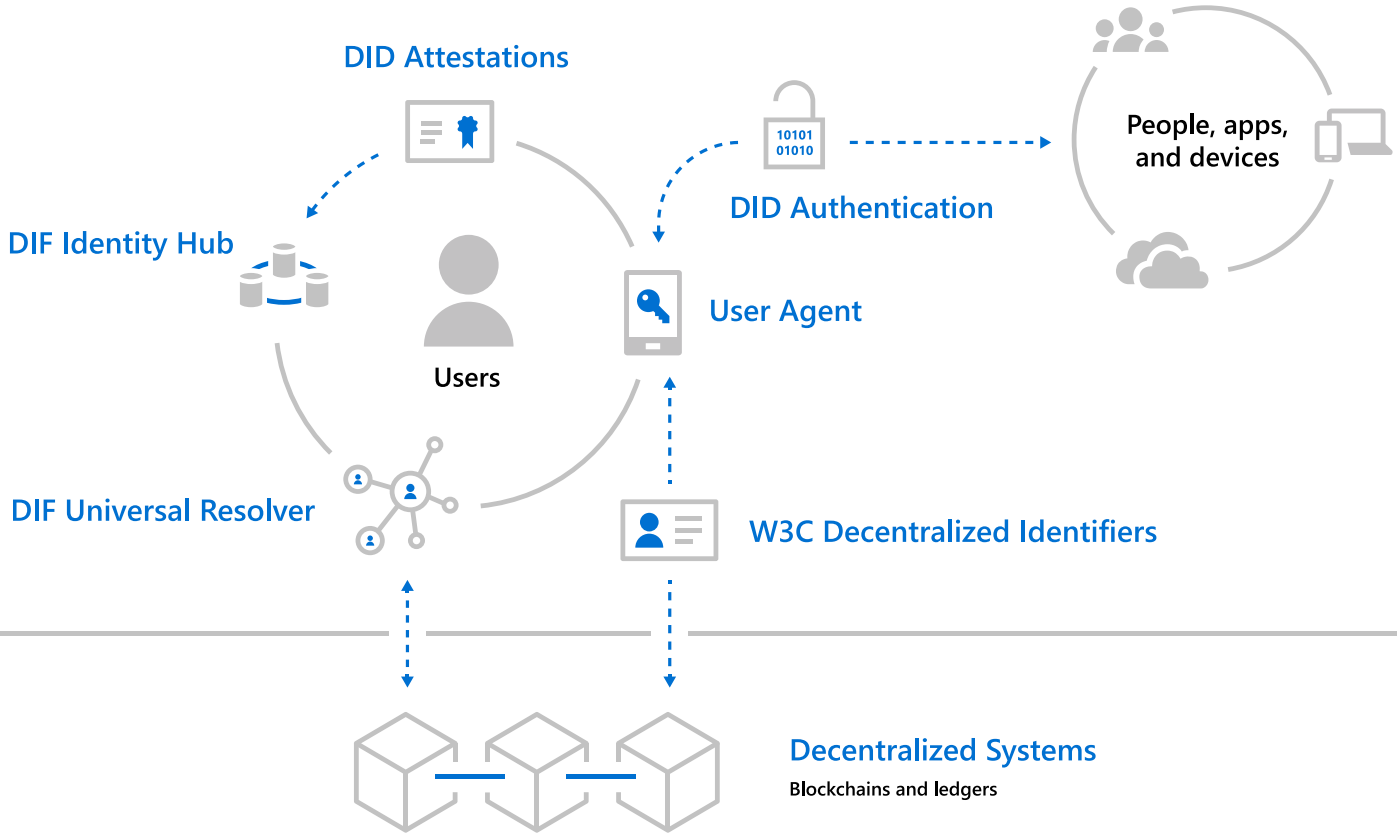

1. W3C Decentralized Identifiers (DIDs) — IDs users create, own, and control independently of any organization or government. DIDs are globally unique identifiers linked to Decentralized Public Key Infrastructure (DPKI) metadata composed of JSON documents that contain public key material, authentication descriptors, and service endpoints.

2. Decentralized systems (for example, blockchains and ledgers) —DIDs are rooted in decentralized systems that provide the mechanism and features required for DPKI. Microsoft is participating in the development of standards and technologies that are being developed by the community to allow for a vibrant ecosystem of DID implementations that support a variety of blockchains and ledgers.

3. DID User Agents —applications that enable real people to use decentralized identities. User Agent apps aid in creating DIDs, managing data and permissions, and signing/validating DID-linked claims. Microsoft will offer a Wallet-like app that can act as a User Agent for managing DIDs and associated data.

4. DIF Universal Resolver—a server that utilizes a collection of DID Drivers to provide a standard means of lookup and resolution for DIDs across implementations and decentralized systems and that returns the DID Document Object (DDO) that encapsulates DPKI metadata associated with a DID.

5. DIF Identity Hubs—a replicated mesh of encrypted personal datastores, composed of cloud and edge instances (like mobile phones, PCs or smart speakers), that facilitate identity data storage and identity interactions.

6. DID Attestations—DID-signed attestations are based on standard formats and protocols. They enable identity owners to generate, present, and verify claims. This forms the basis of trust between users of the systems.

7. Decentralized apps and services—DIDs paired with Identity Hub personal datastores enable the creation of a new class of apps and services. They store data with the user's Identity Hub and operate within the confines of the permissions they are granted.

A sample scenario

Alice has recently graduated from college. She can request a digital copy of her diploma, issued by the university against her DID. She can choose to present her diploma to anyone—like a potential employer—who can independently verify the issuer of the diploma, the time of issuance, and its status.

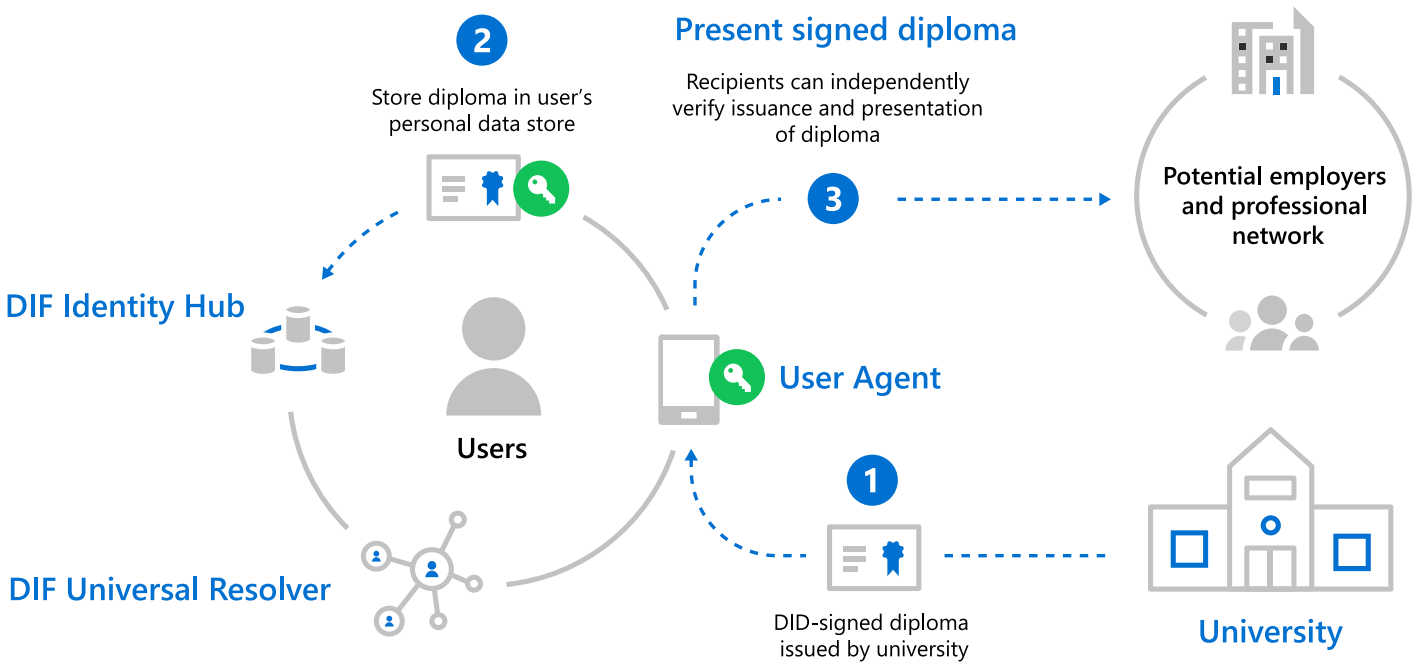

Decentralized Systems

Blockchains and ledgers

Getting started with DIDs

To understand DIDs, it helps to compare them with current identity systems. Email addresses and social network IDs were created as human-friendly aliases for collaboration but are now overloaded to serve as the control points for data access across many scenarios beyond collaboration. This poses a potential problem, given that access to these IDs can be removed at any time by the email provider, social network provider, or other external parties.

Decentralized Identifiers (DIDs) are different. DIDs are user-generated, self-owned, globally unique identifiers rooted in decentralized systems. They possess unique characteristics, like greater assurance of immutability, censorship resistance, and tamper evasiveness. These are critical attributes for any ID system that is intended to provide self-ownership and user control.

Acquiring a DID

To acquire a DID, you use a device under your control to download a DID User Agent app. Just as a web browser is a trusted user agent that helps you navigate the web, a DID User Agent helps you manage all aspects of DIDs—creation of identifiers, authentication, data encryption, and management of keys and permissions. A common misconception about decentralized identity is that all identity data is exposed on public systems like blockchains. This is incorrect. Microsoft believes DID implementations should use decentralized systems strictly to anchor identifiers and non-PII DPKI metadata (as listed above) to enable routing and authentication for the DID owner without risk of censorship. A user's actual identity data resides encrypted "off-chain," under the user's sole control. (For more details, see the "DID interaction using personal datastores" section of this document.)

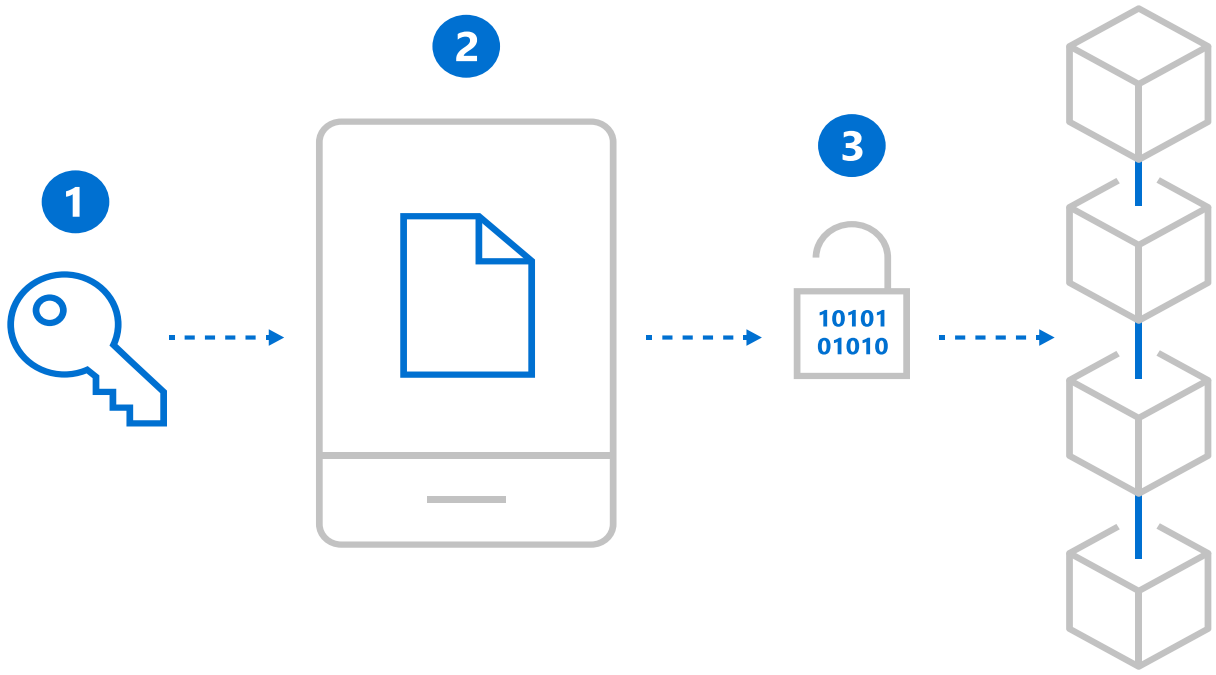

DID interaction using personal datastores

1. DID User Agent app assembles a DID registration payload that includes key references, Identity Hub service endpoints, and public values required for recovery.

2. DID User Agent app generates a device key and an encrypted owner recovery bundle.

3. DID User Agent app pushes the DID registration payload to the decentralized system, in accordance with the protocol of the DID implementation the user selects.

Primary and pairwise DIDs

The broad use of email addresses or social IDs as primary identifiers has led to bad practices. Users have developed a habit of using the same identifier and password across a wide array of products and services. This results in poor security and an association and tracking of accounts.

Conceptually, DIDs can fall into two classes: public DIDs and pairwise DIDs. Public DIDs are IDs that users choose to knowingly link themselves with data intended for the public—for example, a small bio that includes a photograph and a brief description. Public DIDs are suitable if you intend an activity or interaction to be linked to yourself in a way that can be verified by others. But having everything you do tied to a single DID and traceable across the web poses serious privacy and safety risks. This is why pairwise DIDs are useful. Pairwise DIDs are generated whenever users want to isolate their interactions and prevent correlation. For many users, pairwise DIDs might be the primary mechanism they use to conduct identity interactions.

Lookup and discovery of DIDs

Once you have a DID, you might question what you should do if you come across another DID or want to search for one.

DID User Agent apps communicate with DIF Universal Resolver instances to look up DIDs. When a DID is passed to the Universal Resolver, the resolver uses the appropriate driver to interface with decentralized systems and retrieve the matching DID Document.

Some DID implementations are created with the ability to discover all DIDs present in the decentralized system they're rooted in; others lack this capability. Microsoft believes this is a valuable feature because it allows apps and services to deterministically generate a universal directory of DIDs. This is useful in many app and service scenarios that would be difficult without it.

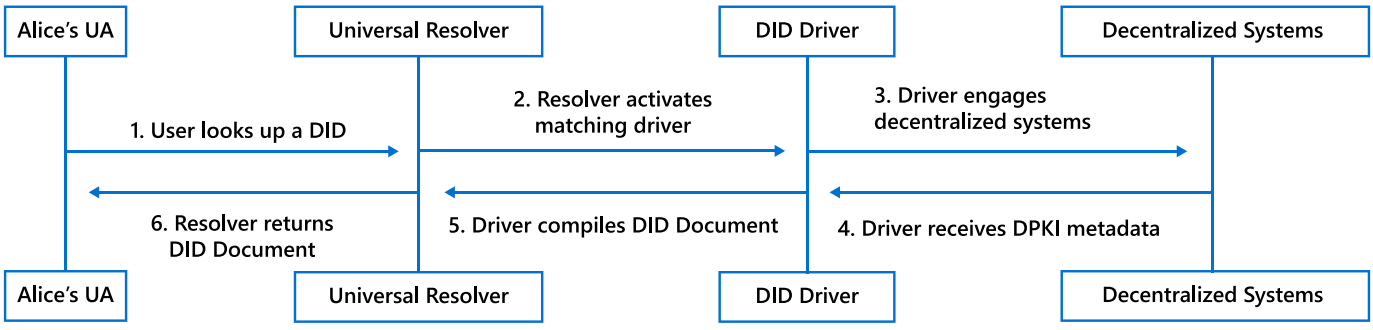

Establishing trust between DIDs

In a world where anyone can create an account or get a DID, how do you know a DID-based identity isn't fake? Much like a personal reputation, DIDs begin life with no evidence of proof; they represent empty identities, and only the owner can prove possession of the DID in question. To accrue evidence of legitimacy, DIDs require endorsements from existing trust providers and processes, like businesses, educational institutions, and governments. DID-based systems provide a mechanism to create attestations that include independent verification of who issued an endorsement and when. By accumulating these attestations from multiple trust systems, an Identity can establish greater confidence over time to match the level of risk inherent in being able to access to an app or service.

As opposed to email identifiers and other current account-based systems, DIDs are self-owned, tied to cryptographic keys, and rooted in decentralized systems that maintain a shared, global lineage of DPKI operations. This enables more advanced identity activities, like the creation and verification of DID-signed attestations. Attestations are independently verifiable claims that one or more DIDs sign with their keys to generate an assertion about another DID. Time-state of data can be logged via blockchain ledgers and independently validated without trusting another entity or organization to record the time of occurrence.

For example, a university might sign an attestation with its DID to substantiate that someone's name is Bob Kelly, that he's a current student, and that his appearance matches his school photo. Universities and other organizations can prove ownership of a DID by demonstrating control of a web domain or through more direct / in-person validation procedures. Testimonies or attestations can then be issued by such organizations against existing credentials and validated with standard cryptographic key suites. The result is that you can require standard, interoperable, verifiable claims from multiple trust providers before engaging in identity interactions and sensitive disclosures to match the level of value being unlocked. This is a game changer for many industries.

Disclosure and authentication of DIDs

There are several ways you can use your DID to interact with another person, app, or service, but the most basic interaction is authentication. To authenticate a DID with an external party, the DID owner discloses a DID to the party. The external party looks up the DID via the DIF Universal Resolver (probably by using a DID User Agent app), and the resolver returns the matching DPKI metadata. The external party generates a challenge by using the public key references in the DPKI metadata and performs a handshake with the user. If the user is able to complete the challenge-response handshake, it's been proven that the user is the owner of DID in question. Microsoft is investigating methods of incorporating support for DID with existing standards (like FIDO/WebAuthn) while maintaining privacy, trust, and security promises.

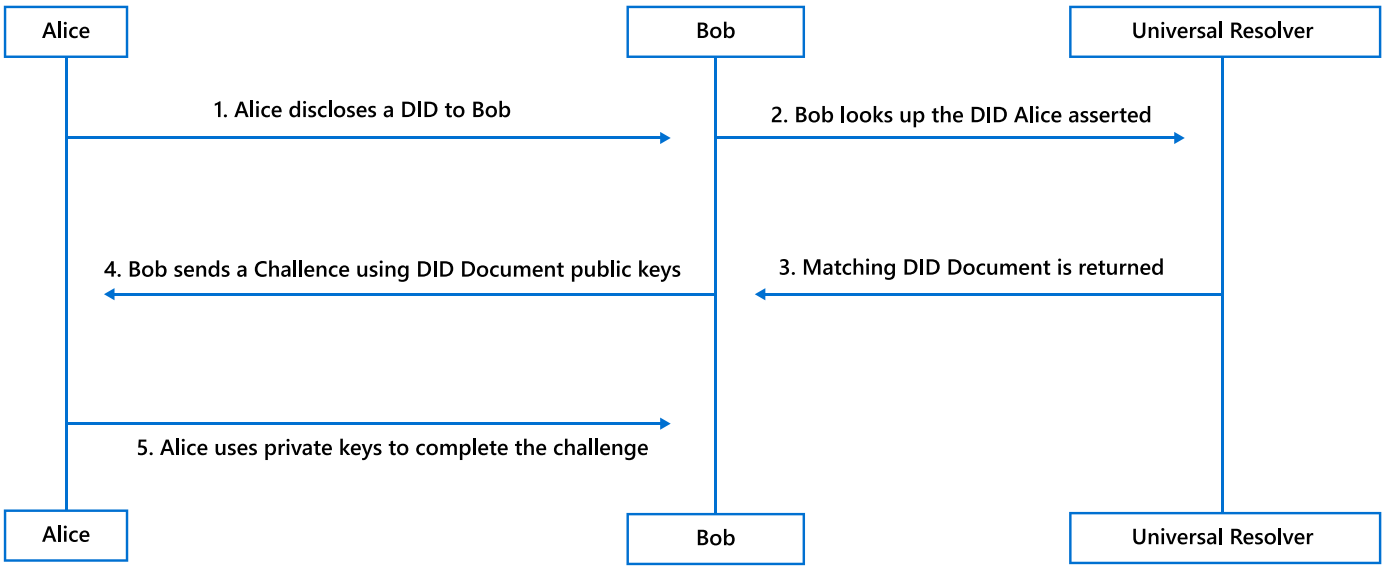

DID interactions using personal datastores

Today, users most commonly store personal data on their local machines or with a provider-based service. While these storage options have their place, they often expose personal data to access by unintended entities and are usually generic storage systems not designed for identity interactions.

DIF Identity Hubs are based on user-controlled, off-chain, personal datastores. Users, via their DID User Agent apps, determine who they want to share data with, and to what level of granularity. Requests to Identity Hubs are routed based on DPKI metadata called Service Endpoints that's associated with DIDs. Identity Hubs are a multi-instance personal mesh, where data is edge-encrypted and user-permissioned to ensure privacy by design. Identity Hubs are designed to support a wide range of identity interactions and provide a foundation for serverless, provider-agnostic, decentralized apps.

Microsoft believes a widely accepted personal datastore standard is the key to unlocking the most compelling use cases in this new ecosystem and is one of many members in DIF working on the DIF Identity Hub specification and reference implementation. Microsoft will offer an instance of DIF's Identity Hub as an Azure service that users can select as one of their Identity Hub instances.

Managing permissions for data access

Identity Hub hosts are facilitators of storage and message relay. They don't have the keys to decrypt user data, and users can revoke/remove encrypted data access from any entity. Permissions are signed with their DID keys, and data is encrypted in accordance with them. Details on the specification for Identity Hub permissions will be made available in the coming weeks.

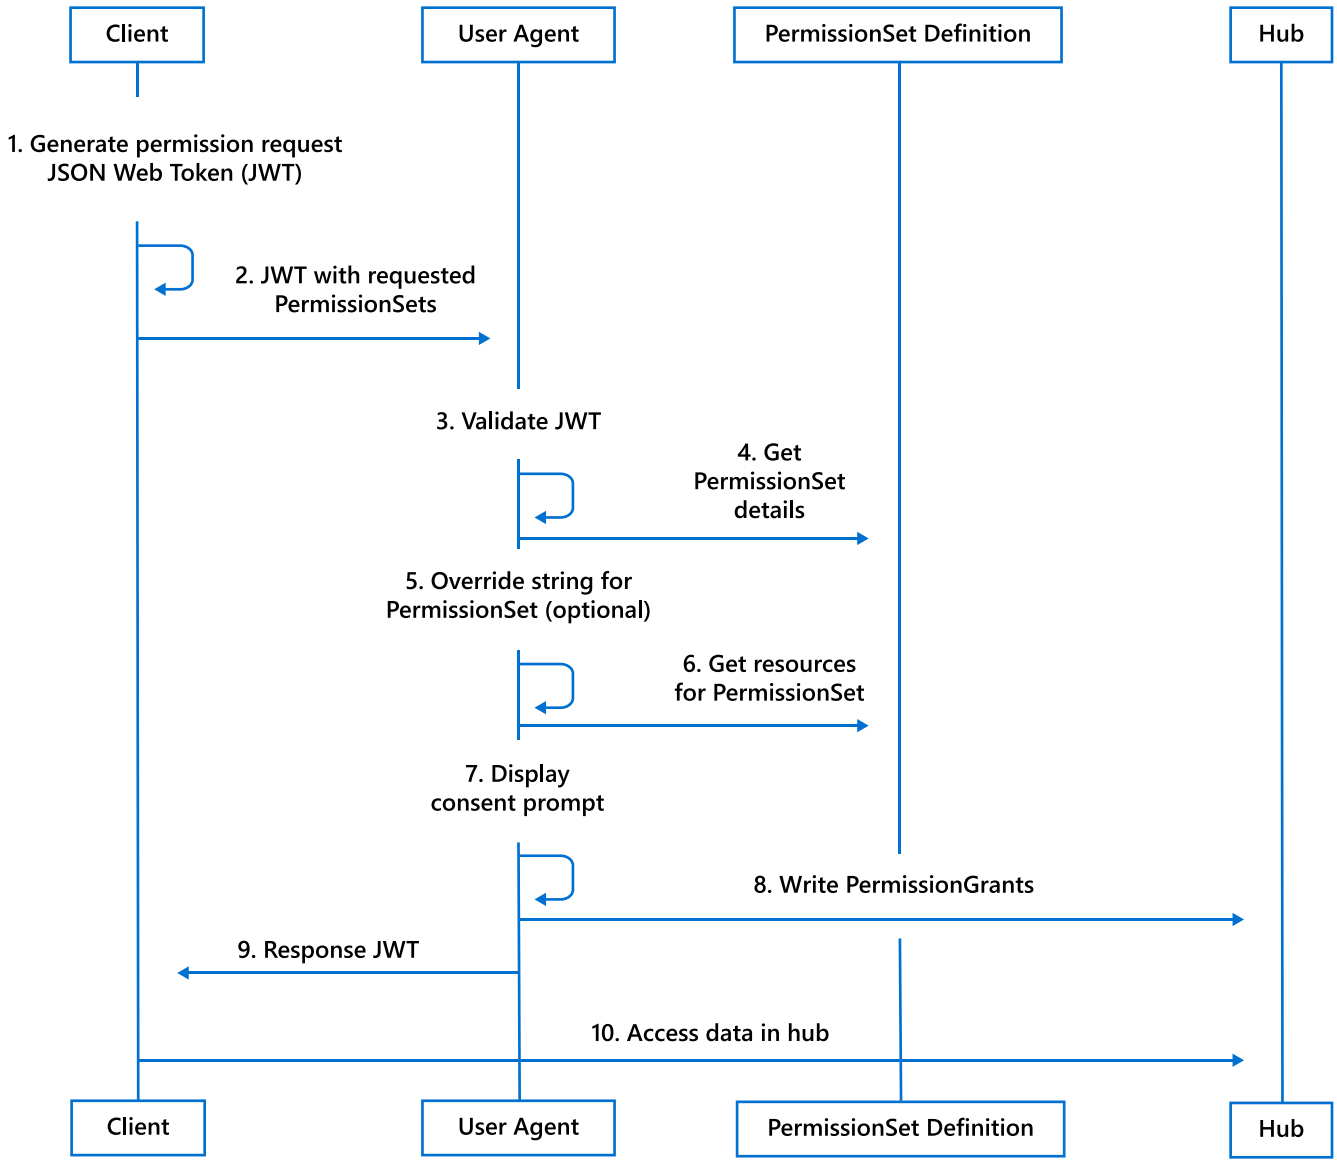

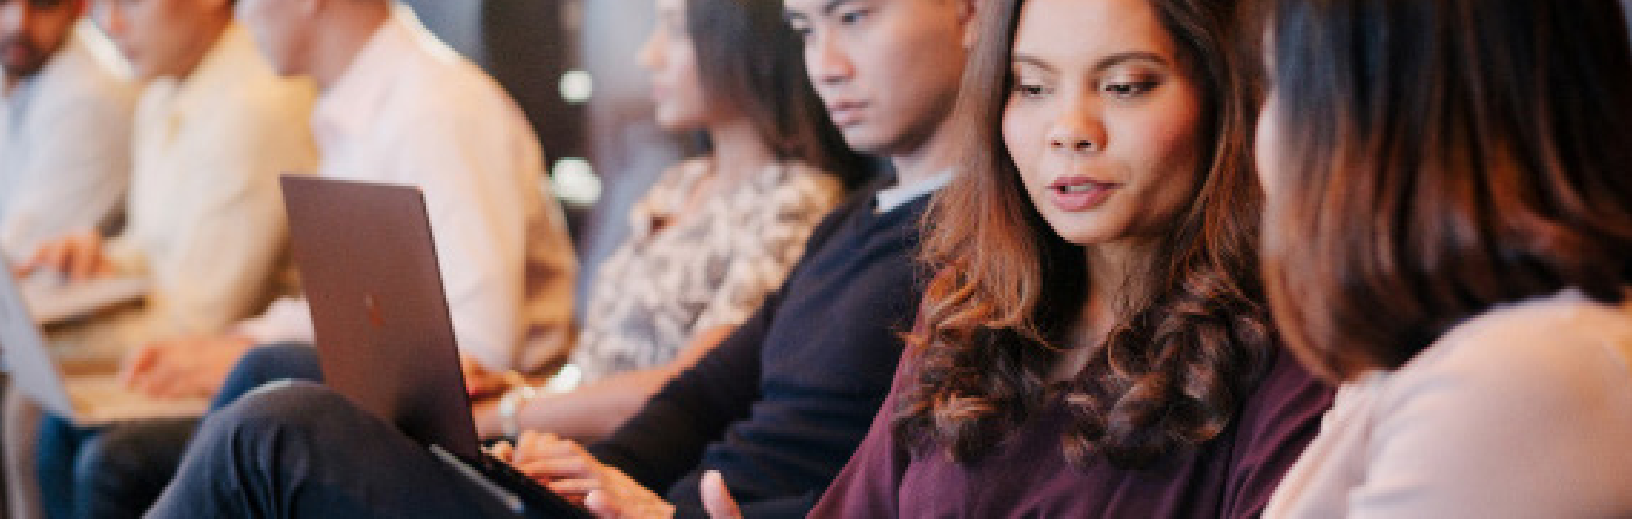

Synchronization and replication of data

A key property of DIF Identity Hubs is that a user can leverage multiple instances across providers and infrastructure boundaries that sync and replicate data to achieve a shared state. But you're not required to use a provider for your Identity Hub at all: Identity Hubs are open source server technology that you can run on any device or infrastructure. This ensures that your identity data is not bound to any organization, upholding the commitment to decentralization, self-ownership, and user control.

Building decentralized apps and services

DIDs paired with Identity Hubs create a foundation of a new type of decentralized application, with features and capabilities that enable a wide variety of uses:

Support for provider-agnostic serverless apps that can be written as purely client-side packages that store their data in the Identity Hubs of users, regardless of where Identity Hub instances are located. This allows apps to interact with any user's Identity Hub through a standard set of APIs that all implementations support.

Semantic, model-less storage of data objects from any shared schema or dataset that apps, services, and organizations use to exchange and collaborate.

An open data layer that anyone can crawl or subscribe to, enabling efficient discovery and awareness of semantic data across all DIDs. DID owners can choose to publish (and when necessary revoke) any type of data, intent, or expression. This creates a vibrant, open marketplace of intended-public data that can be used for P2P offer discovery and value exchange, such as secondhand sales of goods, ride sharing, and vacation rentals.

Here are three examples, across different industries, that illustrate what's possible when data exchange and storage is reoriented in this way:

1. Sellers and gig economy participants can directly publish for-sale items and other offer signals to an open, decentralized market layer by encoding items as semantic Offer objects, and exposing them via Identity Hubs.

2. Patients can allow doctors to interact with their medical data by providing access to their Identity Hub's HL7 FHIR-encoded objects. HL7 FHIR is a schema that's commonly used among doctors, insurance providers, and hospitals.

3. Suppliers and retailers can encode product and service data into their Identity Hubs as GS1 objects, enabling them to exchange supply chain data more efficiently and securely than ever before.

Recovering compromised DIDs

An open question in the DID community is how to handle recovery of DIDs if control is ever compromised via theft or loss of associated keys/devices. We believe that if users are truly the owners of their digital identities, they must be equipped to reliably recover their own DIDs. Microsoft has a hypothesis describing how to empower users to recover their keys on their own. In the coming months, we will share details of our work and code with the community. We look forward to collaborating with you on this important challenge.

Microsoft's progress on Decentralized Identity

Over the past 18 months, Microsoft has invested in incubating a set of ideas for using blockchain and other distributed ledger technologies to create new types of digital identities—identities that are designed from the ground up to enhance personal privacy, security, and control. We aspire to make DIDs a first-class citizen of the Microsoft identity stack.

Over the next few months, we'll provide detailed specs, and, where appropriate, make public code contributions to DID-based technical components, including performance and scale improvements and new tools for DID recovery. Our goal is to help bootstrap this new DID ecosystem by standing up key infrastructure that users and the developer community can depend upon.

Key understandings

1. A user can have one or more DIDs, based on open standards.
2. DIDs can be resolved across chains and ledgers (public, private, and so on).
3. DID permissions are managed via keys accessible only to the user.
4. Identity attributes (or claims) are stored in off-chain DIF Identity Hub personal datastores.
5. Users can have one or more Identity Hub instances, across devices and clouds.
6. User consent is required to access attestations/claims—with granular access controls.
7. Claims are compatible with existing standards (OAuth 2.0 / OIDC).

Get in touch at microsoft.com/ownyouridentity

Join Decentralized Identity Foundation (DIF)
Participate in W3C Credentials Community Group