ReCRED

From Real-world Identities to Privacy-preserving and Attribute-based CREDentials for Device-centric Access Control

Makes your digital life **safe** and definitely **easy**!

www.recred.eu

European Commission | Horizon 2020 European Union funding for Research & Innovation

How Identity federation based on the ID consolidator combined with identity acquisition enhances authentication solutions

In our rapidly evolving online world, every day more and more services tend to shift from offline to online solutions. As the demand and number of web services increases, users need to create multiple online accounts, ending up with multiple online identities one for each service they use. In addition to this, each online service - Service Provider (SP) - maintains its own authentication mechanism with the majority of them using the traditional username/password scheme. This creates the so-called password overload problem where users need to remember one password for each service they maintain.

As a consequence, users resort in re-using the same password for all services, an approach that has serious security vulnerabilities. Second, the user's identity is fragmented across multiple online services (e.g., Facebook, Gmail, Twitter, etc.). This makes the task of proving account joint-ownership of services hard for end-users. Last, there is lack of support for Attribute Based Access Control (ABAC). To clarify, some services require the end-users to authenticate with an account in order to give access to their service revealing their whole identity realm, when they might only need one specific attribute of the end-user's

identity, like age or affiliation. ABAC allows end-users to authenticate to a service by only providing the necessary identity attributes that a specific service is requesting, thus facilitating account-less access by the use of identity attributes, instead of revealing their whole identity.

The aforementioned problems motivate us to design and implement a privacy-preserving architecture for device-centric and attribute-based authentication that aims to anchor all of the users' access control needs to devices (e.g., smartphones) that they habitually carry along. However, the user device becoming the main authentication gateway is not by itself a universal remedy as it entails serious caveats. First, it can become a single point of failure in case of device loss or failure and second it is vulnerable to hijacking after the user has authenticated.

To overcome these problems, we leverage an entity, which we call **Identity Consolidator (IDC)**, in conjunction with behavioral authentication and secure-SIM protocols to efficiently and accurately verify the identity of the user. The IDC enables identity consolidation and management, while at the same time it allows for efficient failure recovery.

Furthermore, our federated architecture allows any SP to delegate the authentication of end-users on a trusted entity called **Identity Provider (IdP)**. SPs are entities that only provide a service to an authorized user, while IdPs are entities that enforce strong authentication mechanisms to securely verify the identity of end-users. In that sense, SPs can adopt our approach with minimal changes to their authentication stacks and/or infrastructure. A great example of a trusted IdP is the IDC. The IDC is the main component in our architecture which comprises the following modules: (a) Online Identity Acquisition module; (b) Physical Identity Acquisition and Verification module; (c) Identity Integration Module; (d) Account Management Module; (e) Identity Management Module; (f) Authentication Management module; and (g) Credential Management module. All these modules interoperate together to solve the problems mentioned above.

First, to address the Identity fragmentation problem, we implement the **Online Identity Acquisition module.** This module is responsible for horizontally binding the online identities (e.g., Facebook account, Twitter account, etc.) of a user by enabling the user to give explicit authorization to ReCRED to access the information of each online account that the user maintains. In addition, some services that require higher attribute assurance level, require the user to enter their physical identity into their system by providing photos of their real-world identities. Since ReCRED aims to fully consolidate the identity of the user, we also implement the **Physical Identity Acquisition module** in order to acquire and verify all the identity attributes included in a user's real-world identity (such as national ID card,

passport, driving license, etc.). The process that an end-user has to follow is separated into two parts, the identity acquisition and the identity verification. The identity acquisition part captures all the physical characteristics and the information included in a user's real-world identity. This can be achieved through NFC-enabled Identity documents or through real-time captured photos of real-world identity documents. Once this information has been voluntarily given to the Physical Identity Acquisition module, it is sent to the Identity Verification part which uses automatic verification techniques (such as OCR, face recognition, and face detection) to verify the validity of the collected identity information. For the verification of this information, the Physical Identity Acquisition module also employs crowdsourcing techniques (peer-to-peer verification of the collected identity information) while ensuring that the privacy of the end-user being verified is guaranteed.

It is worth mentioning that the IDC can know or define the Identity Assurance Level (IAL) of each user identity attribute. The IAL defines how end-users have proven their identity attributes to an identity management system. In the case of the IDC IAL1 is supported through the online registration that the user has to go through when accessing the IDC for the first time. IAL1 requires the identity attributes to be self-asserted. After that, the IAL2 is supported by the Physical and Online Identity Acquisition modules which require remote or in-person identity proofing. Last, the IDC also supports IAL3 which requires in-person identity proofing with the use of Mobile Connect protocol.

Subsequently, the **Identity Repository** is responsible to securely hold all the retrieved and verified identity attributes of the user from both online and physical identity acquisition modules.

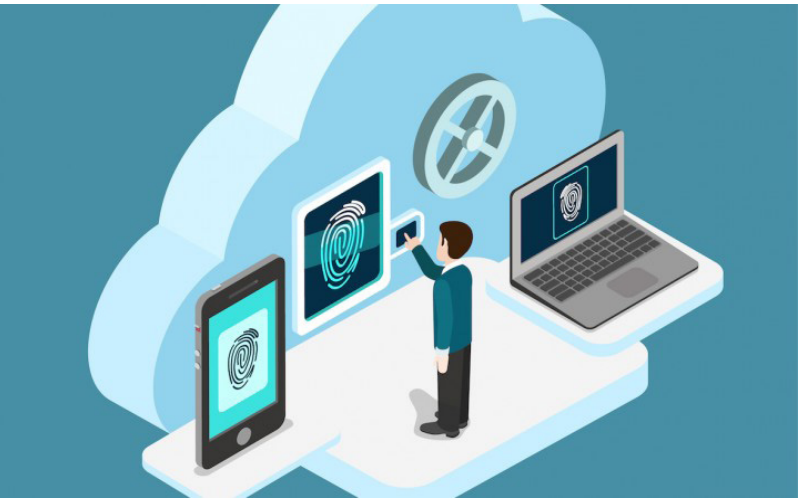

After collecting all this information about an end-user, the IDC needs to normalize the data because some users may have different values for a single identity attribute (e.g., name, birth date, etc.) in different online accounts compared to their real-world identity. ReCRED aims to normalize identity attributes for each user with the use of the **Identity Integration module.** This module receives all the information, if any, gathered by the Online and the Physical Identity Acquisition module and starts aggregating and connecting the acquired attributes, as well as inferring the veracity of the claimed identity attributes. This is achieved using statistical data analysis techniques responsible for standardizing and normalizing this identity information aiming to realize one true value for each identity attribute of the user. The Identity Integration module is also responsible for assigning confidence scores for the veracity of the identity attributes and for labelling them based on their origin. Besides attributes normalization, the Identity Integration module can also infer other or new identity attributes of the user. For example, if the Identity Integration module already knows that a user is a first year student at a

university it may infer that the user is over 18 years old.

Besides providing and verifying her identity attributes, the user needs to also be able manage her account. This is offered by the **Account Management module (AMM)**, which allows the users to manage multiple aspects of their IDC account, e.g., to register IdPs or SPs, or to completely delete their account.

Importantly, through this module, ReCRED aims to protect the accounts of the user in case of device loss or failure, or hijacking. Specifically, the AMM is responsible to keep track of all the Behavioral Authentication Authorities (BAAs). A BAA is a special type of IdP that provides both on-demand and continuous behavioral authentication that can vouch on whether the behavior of a user remains consistent with his usual habits. To achieve this, BAAs continuously track the users' behavior on their devices through various means like keystroke, gait, internet traffic habits, and location to offer a behavioral authentication solution to either SPs or the IDC as a second-factor authentication. In case a behavioral authentication has failed, this may indicates that the device has been stolen, or is no longer held by its legitimate user. As a result, a notification is sent to the AMM to activate LATCH, an account locking functionality. In this case, LATCH locks the user out of all his logged in online services on that specific device. In addition, in case the user lost or had his device stolen, she can activate LATCH manually though the IDC platform.

In general, the AMM is responsible to keep track of all the BAAs, SPs, and IdPs of a user and it also allows BAA, SP, and IdP admins to register their entities with the IDC. Using this knowledge, the AMM can act as a BAA discovery service for the SPs that may require a second-factor authentication. Last, the AMM facilitates the integration of Mobile Connect (Secure-SIM authentication protocol) within our architecture. This means that the user has the ability to verify his mobile phone number through this module. To achieve this, the IDC acts as an SP, requesting Mobile Connect authentication.

After setting up and managing their accounts, the users can use the IDC (as an IdP) to authenticate to any SP that supports ReCRED. The required functionality for authentication is offered by the **Authentication Management Module**. This module implements the functionality required for "killing" the password.

The Authentication Management module offers to the users various

ways to authenticate such as FIDO UAF, Mobile Connect, and behavioral authentication. It extends the OpenID Connect protocol with a FIDO UAF server, thus offering a FIDO-enhanced federated login protocol that allows the IDC to act as an OpenID Connect Provider for undertaking FIDO authentication. The OpenID Connect protocol allows clients to verify and obtain basic profile information about the identity of an end-user based on the authentication performed by an authorization server. OpenID Connect also enables SPs to delegate the authentication of users to IdPs. The FIDO UAF specification, which is a password-less solution, enables the IdPs to authenticate users using strong authenticators, such as biometrics for user to device authentication (e.g., fingerprint) and cryptographic protocols for device-to-service authentication (e.g., RSA). By combining the concepts of strong authentication (FIDO UAF) alongside with the delegation of authentication to IdPs (OpenID Connect) ReCRED allows for a more user-friendly and secure authentication solution for end-users.

In addition, through this module, ReCRED implements a NIST-compliant diverse authentication framework. More precisely, our authentication framework offers various types of authentication methods that determine the granted Authentication Assurance Level (AAL). The AAL defines how end-users can securely authenticate and access an SP. Our architecture targets the highest level of Authentication Assurance (AAL3) which requires hard cryptographic authenticator(s) and two-factor authentication. We support this in our architecture with a Secure-SIM protocol (Mobile Connect), FIDO and IDC backup password (requested during initial registration), or Secure-SIM and FIDO protocols. Here we assume that in the future FIDO and Mobile Connect will be as secure as a hardware cryptographic token (FIPS 140-2) because of advances in TEE and TPMs.

Moreover, each IdP in our architecture has an AAL which denotes the maximum authentication assurance level they can support. IdPs also have an IAL which denotes the maximum identification assurance level they use when they are verifying an identity attribute. As a result, each identity attribute transferred from an IdP inherits the IAL value of the source IdP. Additionally, each identity attribute of a user has both an AAL and an IAL. The AAL of an attribute in the Identity Repository denotes the minimum AAL that the user needs to have in order to access this attribute at a specific IdP. By default, when an attribute

is transferred between IdPs, it inherits the AAL of the source IdP and the destination IdP has to have a policy that gives access to this attribute only if the the user has authenticated with the same AAL or higher.

Apart from these, the Authentication Management module offers the appropriate failure recovery mechanisms that are required to solve the problems that emerge when moving the authentication to the mobile device. The first problem is the single point of failure problem and the second is that in the case the device is stolen the thief has physical access to the device that stores the cryptographic keys. The second problem is addressed via FIDO on devices, which is the human-to-device biometric factor. The first problem is harder to address, yet ReCRED resolves it by leveraging the multiple authentication factors it incorporates.

The IDC federates multiple independent authentication factors offered by Mobile Connect (MC) and BAA to offer a secure and efficient failure recovery mechanism. These independent factors can easily be used in conjunction with a single secure backup password or physical identity verification to reliably authenticate the user during recovery. The user has to first login to the IDC using her secure backup password, which is used only in case of failure recovery. By doing so, the user is granted only temporary and tentative access (AAL1), which provides limited functionality. In particular, the user cannot view, restore or manage credentials and identity attributes. Subsequently, the IDC requires from the user to authenticate with a higher AAL in order to regain full access to her IDC account. In this case the IDC acts as an SP authenticating the user through a Telco IdP via MC. Since the user cannot use FIDO to authenticate, the IDC allows the user to authenticate via SMS using her newly issued by the Mobile Network Operator (MNO) SIM card. In case of device theft or loss and to ensure that the authentication attempt is performed by the legitimate user, the IDC needs to confirm with the MC IdP that the given device was reported as lost and a new SIM card was issued. Additionally, in case the user is not registered with MC then she can use any other OIDC/FIDO-compliant IdP.

For increased assurance and to further verify the identity of the user, the IDC requires from the user to verify her behavior by using one of her trusted BAAs before regaining full access to her IDC account. The user can choose to authenticate to her BAA using a backup password that is specific to the BAA selected by the user during registration with the BAA. Importantly, the user does not have to memorize this BAA backup password since she is able to backup all her BAA backup passwords to the IDC. In addition, BAAs can have insecure and easy to memorize backup passwords as their authentication modality is behavioral and the backup password is used only to prevent denial of service attacks.

After the user has authenticated with the BAA (using the BAA backup password), she is still in tentative access and she is not allowed to manage her behavioral profile. With the user having tentative access, the device sends behavioral records to the BAA, while all the records after the device has been reported as stolen or lost are not considered for the authentication. The IDC acts as an SP while the BAA acts as an IdP authenticating the user based on her behavior. The IDC keeps the user under tentative access until she





successfully proves to the BAA that he/she behaves as she always does. Once the BAA has collected sufficient records to give a verdict on whether the user behaves as usual (that is, similar to how the device was behaving before it was reported lost or was behaviorally detected as stolen) the result is returned to the IDC via OIDC. If the verdict is negative the BAA locks that device out of its IdP. If the verdict is positive, then the user is granted with full access (AAL3) to the IDC and the BAA issues new FIDO credentials for his/her account to the new device. Both MC and BAA authentication is needed because BAA does not increase formally the user's NIST authenticator assurance level but it is just an extra assurance.

If the user does not desire to use any backup password, then she is able to recover from failure with physical identity verification. In this case the user is requested to scan her eID or ePassport using her mobile device. Taking advantage of the NFC capabilities of the device we are able to acquire the verified identity of the user. If the acquired identity matches the one that she had proven to the IDC prior to the failure then she is granted with full access to the IDC.

Based on the NIST standard the user can view and create policies for their identity attributes. This is offered by the **Identity Management Module** (IMM) where users can also view and manage various aspects of their identity. This module empowers users to manage their identity information with the use of the Profile and the Consent Management submodules.

The Profile Management submodule provides easy browsing and management of the identity attributes that IdPs and SPs know about a user and informs the user about the risks of involuntary attributes inference. It also allows users to transfer attribute values between different IdPs by extending federated login protocols like OpenID Connect. This module also allows the IDC to run the federated login protocol OpenID Connect for transferring identity attributes between different IdPs based on the IAL and the AAL of these attributes.

The Consent Management allows users and IdPs to define their consent for their various identity attributes. In this context the users are able to create fine-grained consent policies for their various identity attributes. Here, it is important to note that the NIST IAL policies are supported by our Consent Management Module. In order for users to achieve high levels of identity assurance at a given IdP in most cases they have to release sensitive private information. As a result they may not wish to transfer those attributes of that level

of certainty to another IdP. So, the Consent Management module offers to the user the ability to create policies that prevent the transfer of an attribute with high IAL to an IdP that has an AAL less than the one of the IdP that first acquired that attribute. This means that the user chooses not to store a high certainty attribute to an IdP who has insufficiently secure authentication processes. In addition, this module enables IdP admins to create policies for the identity attributes they maintain. This is required when a source IdP has entailed a significant cost to validate an attribute at a high IAL and may not want to transfer this value to other IdPs. They may instead wish to keep this attribute validity knowledge to themselves as a business competitive advantage. This means that an IdP admin can have policies like "Do not transfer an attribute of IAL3 to another Identity Provider or to a specific Identity Provider.".

Last, ReCRED introduces a simple but privacy-preserving authentication and authorization method. Usually an SP requires from the information such as an email address, phone number, etc., before granting her access to its services. However, this is inconvenient for end-users since it requires to carry or remember this personal information and more importantly, it uniquely identifies the end-user, thus compromising her privacy.

In particular, the IDC offers "Privacy-Preserving Attribute Based Access Control (PABAC)", on top of the OpenID Connect specification, with the **Credential Management module** and taking advantage of the idemix and Uprove Cryptographic credential stacks. PABAC enables SPs that are not aware of our privacy-preserving cryptographic credentials stack to allow end-users to use cryptographic credentials in order to get access to their services and/or resources.

To achieve this, we need a cryptographic credential issuing Idp and a verifying Idp. Users can request the issuance of a cryptographic credential for one or more identity attributes from the issuing Idp, which can be the IDC itself that runs the cryptographic credential issuer stack. The verifying IdP acts as an Idemix/U-Prove verifier able to verify cryptographic credentials while the SPs continue to run the vanilla OpenID Connect protocol.

In other words, the user requests the issuance of a cryptographic credential of one or more user identity attributes from an issuing IdP (e.g., student and over 21). When the user tries to authenticate to an SP, she triggers a session with a cryptographic credential verifying IdP in order for that Verifying IdP to verify the validity of the cryptographic credential. Subsequently, the verifying IdP assures the SP that the user is a holder of a credential that proves that she is a student and is over 21 years old. After this seamless to the user authentication procedure, the user can have access to the service/resources of the SP, knowing that his anonymity is preserved and no more than the needed identity attributes have been revealed to that SP.

Additionally, federated PABAC offers two concepts of anonymity, namely untraceability and unlinkability. This means that that no SP or IdP can track or link any credentials to the user or the other way around. For example, consider the scenario that a user shows a combination of two independent and non-uniquely identifying attributes to an IdP and an SP at time A. If that same user proves the same combination of two cryptographic credentials to another IdP and SP at time B, there is no mechanism that can infer or assure that the user of time A and time B is indeed the same person. This entails untraceability and the users' credentials cannot be linked. This also means that no IdP or SP can build a complete profile of that user since every issuing and verifying session is different.

EXPLOITATION IDifier

In order to exploit some parts of the efforts of the IDC, we have established a spinoff named **IDifier, Ltd**. IDifier offers the Physical Identity Acquisition Module as a service to companies or individuals that require identity verification. We have developed usable web interfaces and a mobile application for soft capturing and verification of real-world identity documents along with the development of the accompanying service infrastructure and algorithms aiming to grow a business.

Online marketplaces and trading forums, the sharing economy, and online social applications are plagued by online fraud rooted in the inability of Internet users to prove their real world identity or parts of it like the age, location, profession, etc. to other users whom they meet online. Motivated by this observation, IDifier's high-level objectives are to:

1. Establish the missing link between a user's online and real-world identity.

2. Offer soft identity acquisition and verification mechanisms to the users allowing them to verify multiple real-world identity documents (such as international IDs, Passports, Driver licenses, etc.) in a fast, reliable and user-friendly manner. IDifier takes advantage of the trusted paths on the devices in order to securely capture images of a user and his real-world identity documents as well as and his physical characteristics (e.g., location). Crowdsourcing techniques and automated means (such as OCR, Facial recognition, Face detection, etc.) are used to verify the acquired information. All the verified identity information is transformed into independent identity attributes which can then be used for the creation of verifiable profiles. Such profiles can be used to prove several identity attributes without revealing your whole identity and while maintaining your anonymity.

3. Offer face-to-face (f2f) verification in which a user physically meets with an

appropriately trained IDifier auditor and presents his documents for validation. F2f audits, especially from highly reputable auditors who can also be eponymous, carry the heaviest weight when it comes to the credibility of a profile.

4. Provide a trustworthy online identity without sacrificing privacy. IDifier allows Internet users to trade and interact with one another as if they had met in person and exchanged credentials like international ID and proof of residence documents. IDifier can be used to send items or money to other people without fear of falling victims of impersonation scams. Additionally, it allows users to chat and/or collaborate online with others that are probably who they claim to be and use a verified identity for safer real-world meetings with people met online.

5. In addition to the above, IDifier can be used in applications that require anonymity, but they need proven identity attributes (e.g., age). For example, IDifier can be used in any online service that requires age verification, allowing a user to prove his age without revealing any other aspect of his identity, enabling him to remain anonymous. At the same time the service can verify the age of the user with high certainty.

6. Using the IDifier's verifiable profiles functionality, a user (e.g., professor) can include a verifiable profile when posting in portals, blogs/forums, or online tabloids allowing co-editors and a reader to verify the identity of the author and also assess his reputation and credibility. In this way, IDifier can contribute in tackling the emerging fake news problem and enables co-editors and readers to easily identify whether an article can be categorized as a fake new and whether it has been posted by a bot or by a credible author.

7. For applications that require anonymity, IDifier provides accountability. For example, a malicious Wikipedia editor that takes advantage of anonymity to vandalize a lemma or spread false information can be held accountable for his actions if he has an IDifier account that bounds his pseudonym with his real world name and his email. IDifier achieves all the above without requiring end-users to jeopardize their privacy by, for example, sending verbatim high quality scans of their ID photo to strangers or revealing their exact address. It provides a systematic solution to all the above problems and replaces the ad hoc, repetitious, and awkward processes used by people today.

A user has to create and validate an IDifier profile only once, and he can subsequently use it multiple times with a simple click, as opposed to having to prove his identity to verifiers with each encounter using questionable and time-consuming ways (e.g., sending scans of passport to strangers and calling them on the phone). IDifier follows security and privacy best practices and its treatment of user's personal information is fully compliant with the EU General Data Protection Regulation (GDPR) (Regulation (EU) 2016/679).


ReCRED


ΠΑΝΕΠΙΣΤΗΜΙΟ ΠΕΙΡΑΙΩΣ
UNIVERSITY OF PIRAEUS


EXUS


institute
iMdea
networks


cnit


Cyprus
University of
Technology


Telefónica
Telefónica I+D


THE
PRODUCTIZERS


certSIGN


wedia


verizon


upcom
Bringing business and IT together


Baker
McKenzie.