

Passwords in Peer-to-Peer

Gunnar Kreitz, Oleksandr Bodriagov, Benjamin Greschbach, Guillermo Rodríguez-Cano, and Sonja Buchegger

KTH Royal Institute of Technology
School of Computer Science and Communication
Stockholm, Sweden
{gkreitz, obo, bgre, gurc, buc}@csc.kth.se

Abstract—One of the differences between typical peer-to-peer (P2P) and client-server systems is the existence of user accounts. While many P2P applications, like public file sharing, are anonymous, more complex services such as decentralized online social networks require user authentication. In these, the common approach to P2P authentication builds on the possession of cryptographic keys. A drawback with that approach is usability when users access the system from multiple devices, an increasingly common scenario.

In this work, we present a scheme to support logins based on users knowing a username-password pair. We use passwords, as they are the most common authentication mechanism in services on the Internet today, ensuring strong user familiarity. In addition to password logins, we also present supporting protocols to provide functionality related to password logins, such as resetting a forgotten password via e-mail or security questions. Together, these allow P2P systems to emulate centralized password logins. The results of our performance evaluation indicate that incurred delays are well within acceptable bounds.

I. INTRODUCTION

Most of the peer-to-peer (P2P) systems deployed today do not authenticate users. While this is often acceptable, or even preferable, there are some problems for which user authentication is a requirement. These include P2P storage, backup, and online social networks. In such applications, the data accessible to a client depends on who is using it.

We discuss how to implement a username-password scheme for authentication in P2P systems. Our goal is to construct an authentication component that can be reused across different P2P applications, which we assume authenticate via possession of cryptographic keys. Thus, from an API perspective, the login system shall allow a user entering a username and a password to recover a set of cryptographic keys which can then be used by the actual application. These keys can also be updated as needed by the application.

The goal from an end-user point of view is to emulate current behavior of centralized password-based login mechanisms. More specifically, we include schemes to remember logins, change passwords, and provide recovery if a password is forgotten. We aim to follow best practice in password authentication, acknowledging that users often re-use passwords between systems. By remembered logins, we mean that a user can opt to have a device store information such that it can log in again without storing the user's password in plain text on the device. Similarly, password change requires knowing the old password, and for password recovery, the user is able to set a new password but does not learn her previous one.

A. Why password authentication?

There is a rich literature on various approaches to authentication, ranging from the traditional username-password pair to hardware tokens and biometry. Of these, the traditional view is that passwords should be replaced by some better mechanism. However, as argued by Herley and van Oorschot [1], despite significant research efforts into dislodging passwords, they are still by far the most common authentication mechanism today. Reasons for their prevalence include simplicity, price, and very strong user familiarity.

When authentication is required in the P2P setting, it is typically done via the security-wise stronger mechanism of generating and storing cryptographic keys on a user's machine. This approach is taken in systems such as OneSwarm [2], Safebook [3], and Tribler [4]. This works well until the user wants to access the service from a second device. To do so, she would need to transfer the keys, or assume a new identity. This is an added complexity and user-perceived drawback for P2P services competing against client-server systems.

One concern is that using passwords may lead to added security risks for skilled and security-conscious users who can easily copy keys between devices. However, nothing prevents such users from choosing passwords of similar strength as cryptographic keys. Another issue pertains to remembered logins, where one must consider theft. We cannot prevent a thief from accessing the user's account, but with our protocols, the thief cannot change the user's password, and the legitimate user can always revoke the remembered credentials that are on the stolen device.

B. Our Contribution

We develop and describe a suite of protocols for password authentication in P2P networks: account registration, login, password change, remembered logins, logout also of a remote device, and password recovery, following best practices and adapting standardized criteria from centralized systems to P2P environments, and start a discussion on usable authentication in P2P systems.

Our password authentication is based on standard cryptographic techniques and can be used with standard P2P components. As a first step toward a security analysis, we discuss the security implications of our protocols. Then, we evaluate the performance of our protocols under various scenarios.

C. Paper Outline

We discuss related work in Section II, give a system overview in Section III and outline our basic scheme for password-based login in Section IV. We then describe password recovery mechanisms as extensions to the basic login mechanisms in Section V. Next, we discuss security in Section VI and report our evaluation results in Section VII before concluding in Section VIII.

II. RELATED WORK

The subject of securely establishing stable identities in P2P systems has been previously studied, for instance by Aberer, Datta and Hauswirth [5]. The need for identities mainly arose from technical concerns, such as handling dynamic IP address assignment, or avoiding Sybil attacks [6]. Authentication of a node is done via a signature key, automatically generated and stored on the node.

As P2P systems began providing more complex functionality [2], [3], [4], [7], the need to authenticate *users*, rather than nodes, arose. It seems that often, authentication via a signature key has been carried over to this problem. While a solution of automatic identification of a node is preferable as long as users use a single device, equating a node with a user fails as users increasingly access services from multiple devices.

Illustrative is the case of backup systems, where an important use case is to restore data on a different system from where it was backed up. Here, two different approaches to authentication have been taken. All approaches build on encrypting backed up content, and the approaches vary in whether the keys are randomly derived [7], or derived from a password [8]. In the former case, a user must manually back the keys up, as these keys are required to restore the backup. The systems deriving a key from a password are related to our proposed protocol, and use some related techniques. However, to the best of our knowledge, they do not consider the additional protocols required surrounding password authentication, such as remembered logins, and recovering lost passwords.

Some P2P storage systems also use techniques which are related to ours. For example, the DHT-based systems GUNet and Freenet use keyword strings to derive a public-private key pair whose private key is used to sign data and the hash of the public key to identify the data in the storage. Both of these systems use a keyword string as a seed to a pseudo-random number generator that produces the key pair [9], [10]. Knowing only the memorable keyword string the user can store and retrieve information.

Related to forgotten passwords, recovery of information in a P2P scenario has been studied by Vu et al. [11] who proposed a combination of threshold-based secret sharing with delegate selection and encrypting shares with passwords.

Frykholm and Juels [12] proposed a password-recovery mechanism based on security questions very similar to our protocol for the same task. They offer better, information-theoretic security properties, something not applicable to our scenario. We treat the subject of password change, which is

not applicable to their scenario, although their proposal could be extended to support password change using our techniques.

III. SYSTEM OVERVIEW AND ASSUMPTIONS

We have designed our system around standard primitives, as depicted in Figure 1. In particular, our protocols build on: a DHT [13], [14], for user lookup; a peer sampling protocol [15], [16] for randomly choosing peers; and a distributed storage [17], [18] for storing data required for our solution. Both the DHT and distributed storage are P2P protocols, run by the peers participating in the system. The storage could be implemented as a DHT, or even be the same as the user lookup DHT. However, we put different requirements on the user lookup DHT and the distributed storage, as detailed below.

To make the system flexible across different implementations, we require as few non-standard features as possible. The exception to this is the DHT that handles account registration, mapping each registered username to a reference in the storage. For resilience against account hijacking, we propose modifying the DHT to be write-once on keys: once an account has been registered, nobody else can register that username.

From the DHT we require two operations, `put(key, value)`, and `get(key)`. The `put` operation associates the *value* with the *key*, and subsequent `get` operations on that *key* will return the *value*. As the DHT is write-once, a second `put` operation with the same *key* will not affect the system state.

The distributed storage functions for data manipulation are similar to the DHT, with three differences: we allow the distributed storage to select the “filename” for us; we require that data can be updated; and we assume (minimalistic) access control when writing. We refer to what is stored in the distributed storage as files, to simplify our description. While the storage component can be implemented as a distributed file system, we emphasize that our requirements are significantly weaker than full file system semantics.

We formalize the API to the storage as having three operations. First, `create(data)` which generates a new file and returns a filename. Second, `write(filename, data)` that overwrites the file *filename* with content *data*. Third, `read(filename)` which reads the content from a file. Our security does not require overwritten data to be inaccessible, so a solution similar to GUNet [10] or Freenet [9] where a new version is stored and pointed to suffice in our protocols.

We require the storage system to support some minimalistic access control. Each stored file has an owner, which is the user who created the file. Only the owner can perform the `write` operation. To authenticate ownership of files, we assume that a public-key cryptographic system is used.

Finally, for the peer sampling component, we require a `getPeer()` method, returning a randomly selected peer, with a distribution close to uniform.

IV. PASSWORD-BASED P2P LOGIN

For password-based authentication in P2P systems, the basic functionality involved is registering an account, and logging in. We also consider password change and remembered logins,




Fig. 1. Overview of the system.

allowing a device to store sufficient information to log in later without asking for credentials anew. Following recommendations from the ISO 27002 standard [19], we define the following requirements for our login procedure and add our own (preceded by a star) to account for several devices.

- passwords should neither be stored nor transmitted in clear text
- a user should be able to choose her own passwords and change them
- files with passwords should be stored separately from application data
- ★ a user should use the same password to log in from any device
- ★ it should not be possible to recover a password by stealing a device with remembered credentials
- ★ it should be possible to block access to the account from a stolen device

The standard also defines limitations for password login procedures that our system cannot provide fully due to the lack of rate-limiting possibilities in P2P networks: to limit the number of unsuccessful login attempts and the maximum and minimum time allowed for the login procedure. Adapting a multi-party password hardening scheme [20] could, in future work, be a way to achieve similar properties in a P2P network. Besides this limitation, our protocols fulfill the requirements as outlined in the standard, and our own added requirements.

We now describe our protocols based on the system model from Section III. Figure 2 shows the information objects and their storage locations, with arrows for the abstract flow of the login procedure, Table I lists the terms used in the algorithms.

A. Account Registration

To register a new account (see Algorithm 1), the user first has to choose a username $uname$ and a password $passwd$. Next, the user creates a key store file F_{KS} , containing all the keys used by the P2P application the user wants to log in to

TABLE I
PROTOCOL TERMINOLOGY

$uname$	Username
$passwd$	Password
$salt$	Random byte string
K_W	Cryptographic key for write authentication
F_{KS}	Key store file
f_{KS}	File name of F_{KS}
K_{KS}	Cryptographic key (used to encrypt F_{KS})
F_{LI}, f_{LI}, K_{LI}	Login information file, its file name and key
F_{DL}, f_{DL}, K_{DL}	Device login information file, its file name and key
D, D_{ID}	User device and the identifier of D
K_{x1}, K_{x2}, \dots	Cryptographic keys for usage after logging in
$devmap$	Mapping from device identifiers to device login information files and corresponding keys

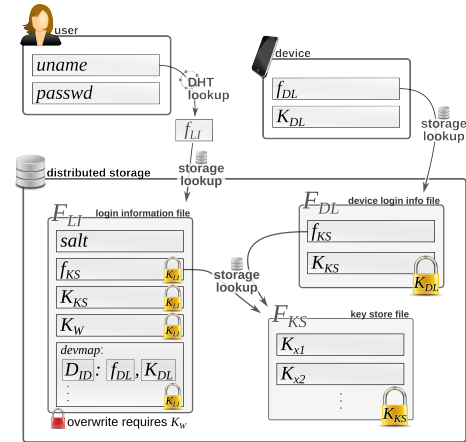


Fig. 2. Storage Locations (boxes) and Login Procedure (arrows)

(and an additional storage key, authenticating write operations on this file). The user then creates a symmetric key K_{KS} , encrypts the file content with this key and puts the ciphertext into the storage, obtaining a file name f_{KS} . Now, the user creates a login information file F_{LI} by creating a random byte string $salt$, deriving a symmetric key K_{LI} from the password $passwd$ and the $salt$, encrypting f_{KS} , K_{KS} and K_W (a generated storage key, required for overwriting F_{LI} later) with K_{LI} . The salt and the three encrypted values are put into the storage, obtaining a file name f_{LI} . The salt is stored in plaintext, so that the user later can derive the decryption key K_{LI} by only providing the password. Finally, the user performs the write-once operation put on the DHT with $uname$ as key and f_{LI} as value. If the username was taken, the user is prompted for a new username. Once all operations have succeeded, the user is registered in the system.

B. Login

Once registered, a user is able to log in – that is, to retrieve the cryptographic keys stored in the key store file F_{KS} – from any device by only entering her username and password (see Algorithm 2). A get request with the parameter $uname$ to the DHT results in the filename f_{LI} for the login information file F_{LI} . This file is retrieved from the distributed storage and contains the $salt$ in plaintext. The latter is fed into a key-

Algorithm 1 Account Registration

```
1:  $uname \leftarrow$  User.input("Choose username:")
2:  $passwd \leftarrow$  User.input("Choose strong password:")
3:  $K_{KS} \leftarrow$  generateKey()
4:  $F_{KS} \leftarrow$  encrypt $_{K_{KS}}$ ( $K_{x1} || K_{x2} || \dots$ )
5:  $f_{KS} \leftarrow$  Storage.create( $F_{KS}$ )
6:  $salt \leftarrow$  generateSalt()
7:  $devmap \leftarrow$  createMap()
8:  $K_{LI} \leftarrow$  KDF( $salt, passwd$ )
9:  $K_W \leftarrow$  generateKey() // suitable for the storage system
10:  $F_{LI} \leftarrow salt || encrypt_{K_{LI}}(f_{KS} || K_{KS} || K_W || devmap)$ 
11:  $f_{LI} \leftarrow$  Storage.create( $F_{LI}$ ) // using  $K_W$ 
12: while DHT.put( $uname, f_{LI}$ ) fails
13:    $uname \leftarrow$  User.input("Choose new username:")
14: end while
```

derivation function together with the user password to derive the key K_{LI} . This key allows the user to decrypt all other content of the login information file, including the filename f_{KS} of the key store file and the corresponding key K_{KS} . Finally, the user fetches the key store file F_{KS} from the storage system and decrypts it, using K_{KS} . This concludes the login procedure as the user is now in possession of the keys K_{x1}, K_{x2}, \dots , required by the P2P application.

If the user chose to remember the login information on the local device, a new device login information file F_{DL} is created and saved to the storage system (which returns a filename f_{DL}). This file contains the filename f_{KS} of the key store file as well as the according key K_{KS} and is encrypted with a new key K_{DL} . On the device, only the filename f_{DL} and the key K_{DL} are stored locally. Additionally, a reference to the device login information file is stored in the $devmap$ value of the login information file F_{LI} . It contains a mapping from a device identifier to the filename and key of the device login information file, allowing password changes and device revocation as described later.

When the user wants to log in from the same device again, the locally stored values (f_{DL}, K_{DL}) are used to retrieve the device login information file, decrypt it, and thereby gain access to the key store file. Thus, the remembered login feature allows the user to log in without entering the password, while nothing password-related is stored on the device. Furthermore, remembered logins remain valid even if the P2P application changes keys in the key store file.

C. Password Change

Before the user can change the password, she must log in using her password to obtain K_{LI} . With this information, the password change can be accomplished (see Algorithm 3): the user is asked for a new password and a new salt is generated. The key-derivation function is used to generate a new key K_{LI}^{new} for the login information file. Then, the content of the key-store file is fetched and decrypted (with the old key). A new key K_{KS}^{new} is generated and used for encrypting the key-store content again before it is saved to the storage

Algorithm 2 Login

```
1:  $f_{DL}, K_{DL} \leftarrow$  Device.readLocalStore()
2: if  $f_{DL} \neq$  NULL then // non-interactive login
3:    $F_{DL} \leftarrow$  Storage.read( $f_{DL}$ )
4:    $f_{KS}, K_{KS} \leftarrow$  decrypt $_{K_{DL}}$ ( $F_{DL}$ )
5:    $saveLoginLocally \leftarrow$  False
6: else // interactive login
7:    $uname \leftarrow$  User.input("Enter username:")
8:    $passwd \leftarrow$  User.input("Enter password:")
9:    $saveLoginLocally \leftarrow$  User.input("Remember?")
10:   $f_{LI} \leftarrow$  DHT.get( $uname$ )
11:   $F_{LI} \leftarrow$  Storage.read( $f_{LI}$ )
12:   $salt \leftarrow F_{LI}.salt$  // stored in plaintext
13:   $K_{LI} \leftarrow$  KDF( $salt, passwd$ )
14:   $f_{KS}, K_{KS}, K_W, devmap \leftarrow$  decrypt $_{K_{LI}}$ ( $F_{LI}$ )
15: end if
16:  $F_{KS} \leftarrow$  Storage.read( $f_{KS}$ )
17:  $K_{x1}, K_{x2}, \dots \leftarrow$  decrypt $_{K_{KS}}$ ( $F_{KS}$ )
18: if  $saveLoginLocally$  then
19:    $K_{DL} \leftarrow$  generateKey()
20:    $F_{DL} \leftarrow$  encrypt $_{K_{DL}}$ ( $f_{KS} || K_{KS}$ )
21:    $f_{DL} \leftarrow$  Storage.create( $F_{DL}$ )
22:   Device.writeLocalStore( $f_{DL} || K_{DL}$ )
23:    $devmap.append$ (Device.ID,  $f_{DL} || K_{DL}$ )
24:    $F_{LI} \leftarrow salt || encrypt_{K_{LI}}(f_{KS} || K_{KS} || K_W || devmap)$ 
25:   Storage.write( $f_{LI}, F_{LI}$ ) // using  $K_W$ 
26: end if
```

system, obtaining a new filename f_{KS}^{new} . Finally, the login information file is updated: $f_{KS}^{new}, K_{KS}^{new}$, the write credential K_W as well as a new empty device mapping $devmap^{new}$ are encrypted with the new key K_{LI}^{new} . Together with the new salt, this ciphertext is written to the distributed storage, using the reference f_{LI} and the credential K_W , to authenticate the write operation. Lastly, the keys stored in the key store should be updated by the application using our P2P protocol. See Section VI-E for a discussion. At this point, old device login information files can also be deleted from the storage to reclaim space.

D. Logout

To log out from the system, the user does not have to interact with the DHT or the storage system. Simply wiping her local cache from application data and all key material restores the pre-login state. If the user chose to remember the login on a device, the corresponding device login information file F_{DL} can also be deleted from the storage.

A problem related to logging out is revoking remembered credentials on another device, e.g., a user's stolen phone. To accomplish this, we first run the password change operation, which locks out all devices with remembered logins, because the key store key K_{KS} changed (as well as the filename f_{KS}). Next, we use the device mapping $devmap$ to inform all devices about the new key (and filename), except the device that is to be revoked. To inform a device about the change, we update

Algorithm 3 Password Change

Input: $uname, K_{LI}^{old}$

- 1: $f_{LI} \leftarrow \text{DHT.get}(uname)$
- 2: $F_{LI}^{old} \leftarrow \text{Storage.read}(f_{LI})$
- 3: $f_{KS}^{old}, K_{KS}^{old}, K_W, devmap^{old} \leftarrow \text{decrypt}_{K_{LI}^{old}}(F_{LI}^{old})$
- 4: $passwd^{new} \leftarrow \text{User.input}(\text{"Enter new password:"})$
- 5: $salt^{new} \leftarrow \text{generateSalt}()$
- 6: $K_{LI}^{new} \leftarrow \text{KDF}(salt^{new}, passwd^{new})$
- 7: $devmap^{new} \leftarrow \text{createMap}()$
- 8: $F_{KS}^{enc-old} \leftarrow \text{Storage.read}(f_{KS}^{old})$
- 9: $F_{KS} \leftarrow \text{decrypt}_{K_{KS}^{old}}(F_{KS}^{enc-old})$
- 10: $K_{KS}^{new} \leftarrow \text{generateKey}()$
- 11: $F_{KS}^{enc-new} \leftarrow \text{encrypt}_{K_{KS}^{new}}(F_{KS})$
- 12: $f_{KS}^{new} \leftarrow \text{Storage.create}(F_{KS}^{enc-new})$
- 13: $F_{LI}^{new} \leftarrow salt^{new} || \text{encrypt}_{K_{LI}^{new}}(f_{KS}^{new} || K_{KS}^{new} || K_W || devmap^{new})$
- 14: $\text{Storage.write}(f_{LI}, F_{LI}^{new})$ // using K_W
- 15: Refresh keys stored in key store
- 16: Old device login information files may be deleted

the corresponding values in the device's login information file F_{DL} which can be accessed from the device by using the locally stored credentials.

Algorithm 4 describes this necessary extension. After running the password change operation, all devices that should *not* be revoked and that have remembered logins (and therefore are referenced in the device mapping $devmap$) are processed. The device login information filename f_{DL} and its key K_{DL} are read, and the new key store key K_{KS}^{new} and filename f_{KS}^{new} are written to the device login information file F_{DL} , encrypted under the device key K_{DL} . Finally, the modified $devmap$ is saved back to the login information file F_{LI} .

Algorithm 4 Logout Other Device

- 1: ... // run Algorithm 3 (Password Change)
- 2: $deviceToLogout \leftarrow \text{User.input}(\text{"Select device:"})$
- 3: $devmap.remove(deviceToLogout)$
- 4: **foreach** D_{ID} **in** $devmap$ // all devices to keep
- 5: $f_{DL}, K_{DL} \leftarrow devmap.get(D_{ID})$
- 6: $F_{DL} \leftarrow \text{encrypt}_{K_{DL}}(f_{KS}^{new} || K_{KS}^{new})$
- 7: $\text{Storage.write}(f_{DL}, F_{DL})$
- 8: **end**
- 9: ... // save modified $devmap$ back to F_{LI}

V. PASSWORD RECOVERY

An important part of password-based logins is the possibility for users to recover their accounts if they forget their passwords. We refer to this as a *password recovery mechanism*. The goal of a password recovery mechanism is to provide a secondary way of authenticating the user. There are a number of password recovery mechanisms used in practice. In our experience, three of the most common ones are password hints, security questions, and e-mail based recovery. Other

approaches (beyond the scope of this paper) include vouching for identity by social contacts [21], or using trusted devices.

Password hints means that the user may enter a hint at the same time as she sets this password. The hint will be displayed to her if she forgets her password, and should be selected such that it helps her recall her password, but does not make it significantly easier for someone else to guess it. The hint is not truly a secondary authentication mechanism, but rather a means to recovering the original password-based authentication mechanism. A basic version of password hints would be straightforward to implement in our system: the hint can be stored in plaintext in the login information file. Security questions and e-mail based password recovery are more complex to adapt. We described their implementation in detail after listing requirements.

As in Section IV for the login procedure, we define a set of functional requirements for password recovery, based on the ISO 27002 standard [19] as follows. We also augment the list with requirements of our own (preceded by a star).

- establish methods to verify the identity of a user prior to allowing the user to choose a new password
- communicate with those affected by or involved with recovery security incidents
- have procedures to allow recovery and restoration of business operations and availability of information in a time-scaled manner
- a legitimate user should be able to recover lost (forgotten) or broken (device's) keys
 - ★ the recovery procedure should allow a user to set a new password, not reveal the old password
 - ★ the process of recovery should be easy to use
 - ★ sensitive information for recovery should be kept secret

Our protocols support these requirements. The sole exception is that if a password is reset via security questions alone, the system would not "communicate with those affected" (e.g., send an e-mail notification that the password had been reset, as is common in centralized services). We remark that the last item is a stronger property than many centralized systems provide. In our system, no one learns the answers to a user's security questions. We consider this to be important, since many systems use similar security questions.

The operations described in this section imply minor additions to the protocols of Section IV, i. e., invoking the update procedures after each password change (to sustain transaction safety, the updates have to be included in the final write operation of the password change operation).

A. Security Questions

Security questions is a password recovery technique that relies on answers to questions the user is asked during registration. The answers should be such that they cannot be easily guessed or researched by an attacker, but still stable over time, memorable, and definite [22]. Rabkin [23] underlines the importance to choose good questions especially in the era of social networks. Frykholm and Juels [12] discuss a related technique that is similar to our adaption of this scheme.

TABLE II
RECOVERY PROTOCOL TERMINOLOGY

Security Question Recovery	
qS_i	(n, k) -secret sharing share of K_{LI}
Q_i	Security challenge question
A_i	Answer to question Q_i
$qsalt_i$	Random byte string
qK_i	Key to encrypt the share qS_i
E-mail Based Recovery	
K_R	Long-term recovery key
eS_i	(n, k) -secret sharing share of K_R
$email$	Recovery e-mail address of the user
$peer_i$	Randomly selected peer
$esalt_i$	Random byte string (to seed the e-mail commitment)
$ksalt_i$	Random byte string (to seed the key eK_i)
C_i	Cryptographic commitment to the e-mail address
eK_i	Key to encrypt the share eS_i

We assume that the user provides n answers A_i to suitable security questions Q_i . In order to recover the password, we require the user to answer any k out of these n questions correctly. The choice of k constitutes an obvious trade-off between security and usability. A successful recovery yields the key K_{LI} to the login information file, allowing the user to change the password, using Algorithm 3. Our implementation does not require the user to provide new answers after a regular password change. Additionally, we avoid storing the plaintext answers to the security questions.

For the setup of the question based recovery mechanism (Algorithm 5), we first create n shares qS_1, \dots, qS_n of the key K_{LI} under an (n, k) -secret sharing scheme. For each of these shares, we create a salt $qsalt_i$, derive a key qK_i from this salt and the answer A_i , and use it to encrypt the share, yielding qS_i^{enc} . Furthermore we encrypt the key qK_i with the login information file key K_{LI} , for the update procedure described later. Finally, the login information file is extended with all questions Q_i , the salts $qsalt_i$, the encrypted shares qS_i^{enc} and the encrypted keys qK_i^{enc} . When recovering, the user has to reproduce at least k answers, which together with the stored salts can be used to derive k keys qK_i , which in turn can decrypt k shares qS_i .

When K_{LI} changes (e.g., due to a regular password change), we update the recovery information as in Algorithm 6: for the new key K_{LI}^{new} , a new set of shares is created. Next, the keys qK_i are decrypted and used to encrypt the new shares. Neither the keys qK_i nor the salts $salt_i$ change, so the user can still use the same answers for recovery. Finally, the updated shares (and re-encrypted keys, to allow further updates) are saved back to the login information file.

B. E-mail Based

In e-mail based password recovery, the user is sent an e-mail containing some information, typically a link with a token, by which she can reset her password. This link is sent to an e-mail address she has registered with her account.

We adapt this scheme by randomly choosing a number of peers, that collaboratively provide this functionality to the user.

Algorithm 5 Security Questions Setup

```

1:  $qS_1, \dots, qS_n \leftarrow \text{createShares}(n, k, K_{LI})$ 
2: for  $i \leftarrow 1, n$  do
3:    $Q_i \leftarrow \text{User.input}(\text{"Enter question } i:\text{"})$ 
4:    $A_i \leftarrow \text{User.input}(\text{"Enter answer } i:\text{"})$ 
5:    $qsalt_i \leftarrow \text{generateSalt}()$ 
6:    $qK_i \leftarrow \text{KDF}(qsalt_i, A_i)$ 
7:    $qS_i^{enc} \leftarrow \text{encrypt}_{qK_i}(qS_i)$ 
8:    $qK_i^{enc} \leftarrow \text{encrypt}_{K_{LI}}(qK_i)$ 
9: end for
10: add to  $F_{LI}$ :  $qS_i^{enc}, qK_i^{enc}$  and the plaintext values of
     $Q_i, qsalt_i \quad \forall i \in \{1, \dots, n\}$ 

```

Algorithm 6 Security Questions Update (on K_{LI} change)

```

1:  $qS_1^{new}, \dots, qS_n^{new} \leftarrow \text{createShares}(n, k, K_{LI}^{new})$ 
2: for  $i \leftarrow 1, n$  do
3:    $qK_i \leftarrow \text{decrypt}_{K_{LI}}(qK_i^{enc})$ 
4:    $qS_i^{new-enc} \leftarrow \text{encrypt}_{qK_i}(qS_i^{new})$ 
5:    $qK_i^{new-enc} \leftarrow \text{encrypt}_{K_{LI}^{new}}(qK_i)$ 
6: end for
7: update in  $F_{LI}$ :  $qS_i^{new-enc}, qK_i^{new-enc} \quad \forall i \in \{1, \dots, n\}$ 

```

We use (n, k) -secret sharing to enable password recovery even if not all of the involved peers are online when the user wants to recover the password. We discuss parameter choices of k and n in Section VII-B.

To provide persistence of the recovery mechanism independent of a changing key K_{LI} (e.g., due to a password change), the result of the recovery process is a recovery key K_R , that always encrypts the current version of K_{LI} . Algorithm 7 describes the setup procedure: From the recovery key K_R , n shares eS_1, \dots, eS_n are generated using (n, k) -secret sharing. For each share, a random peer $peer_i$ is picked, two salts $esalt_i$ and $ksalt_i$ are created and a cryptographic commitment C_i is derived from the salt $esalt_i$ together with the $email$. This commitment will be used to authorize the user to the peer, and bind it to this specific e-mail address. Next, a key eK_i , to encrypt the share eS_i , is derived in the same way as the commitment, but with salt $ksalt_i$. A different salt is needed so that the peer cannot decrypt the share (before learning the address). The commitment and the encrypted share are stored at the peer. The login information F_{LI} file is extended with a list of the chosen peers $peer_i$ and the according salts $esalt_i, ksalt_i$, as well as K_{LI}^{enc} , encrypted with the recovery key, and the recovery key, encrypted with K_{LI} (to allow for password changes).

To recover the password (Algorithm 8), the user looks up the available information in the login information file, including the list of peers to be requested for assistance. Each request is authorized by the commitment C_i , that the peer can derive from the salt $esalt_i$ and the e-mail address, that the user provided (Algorithm 9). If the request was legitimate, the peer sends the encrypted share to the e-mail address. As soon as the user collected k answers, she can recover K_{LI} .

Algorithm 7 E-mail Recovery Setup

```
1:  $K_R \leftarrow \text{generateKey}()$  // long-term recovery key
2:  $K_{LI}^{enc} \leftarrow \text{encrypt}_{K_R}(K_{LI})$ 
3:  $eS_1, \dots, eS_n \leftarrow \text{createShares}(n, k, K_R)$ 
4:  $email \leftarrow \text{User.input}(\text{"Enter recovery e-mail address:"})$ 
5: for  $i \leftarrow 1, n$  do
6:    $peer_i \leftarrow \text{getPeer}()$ 
7:    $esalt_i \leftarrow \text{generateSalt}()$ 
8:    $ksalt_i \leftarrow \text{generateSalt}()$ 
9:    $C_i \leftarrow \text{KDF}(esalt_i, email)$  // commitment
10:   $eK_i \leftarrow \text{KDF}(ksalt_i, email)$ 
11:   $eS_i^{enc} \leftarrow \text{encrypt}_{eK_i}(eS_i)$ 
12:  store at  $peer_i$ :  $C_i, eS_i^{enc}$ 
13: end for
14:  $K_R^{enc} \leftarrow \text{encrypt}_{K_{LI}}(K_R)$ 
15: add to  $F_{LI}$ :  $K_{LI}^{enc}, K_R^{enc}$  and the plaintext values of
     $peer_i, esalt_i, ksalt_i \quad \forall i \in \{1, \dots, n\}$ 
```

Algorithm 8 E-mail Recovery: User

```
1:  $uname \leftarrow \text{User.input}(\text{"Enter username:"})$ 
2:  $email \leftarrow \text{User.input}(\text{"Enter e-mail:"})$ 
3:  $f_{LI} \leftarrow \text{DHT.get}(uname)$ 
4:  $F_{LI} \leftarrow \text{Storage.read}(f_{LI})$ 
5:  $K_{LI}^{enc}, \forall i: peer_i, esalt_i, ksalt_i \leftarrow F_{LI}$  // plaintext part
6:  $\forall i$ : send  $(email, esalt_i)$  to  $peer_i$  // send  $n$  requests
7:  $eS_1^{enc}, \dots, eS_k^{enc} \leftarrow \text{read e-mail}$  // wait for  $k$  e-mails
8: for  $i \leftarrow 1, k$  do
9:    $eK_i \leftarrow \text{KDF}(ksalt_i, email)$ 
10:   $eS_i \leftarrow \text{decrypt}_{eK_i}(eS_i^{enc})$ 
11: end for
12:  $K_R \leftarrow \text{useShares}(eS_1, \dots, eS_k)$ 
13:  $K_{LI} \leftarrow \text{decrypt}_{K_R}(K_{LI}^{enc})$ 
14: ... // run Algorithm 3 (Password Change)
```

To provide long-term persistence of this recovery mechanism, K_{LI}^{enc} has to be updated whenever K_{LI} changes. Algorithm 10 describes the necessary steps, including updating K_R^{enc} to allow subsequent updates.

Algorithm 9 E-mail Recovery: Peer

```
Stored:  $C_i, eS_i^{enc}$  // stored at peer
Input:  $email, esalt_i$  // provided by the user request
1: if  $C_i = \text{KDF}(esalt_i, email)$  then // legitimate request
2:   sendMail( $email, eS_i^{enc}$ )
3: end if
```

Algorithm 10 E-mail Recovery Update (on K_{LI} change)

```
1:  $K_R \leftarrow \text{decrypt}_{K_{LI}^{old}}(K_R^{enc})$ 
2:  $K_{LI}^{enc-new} \leftarrow \text{encrypt}_{K_R}(K_{LI}^{new})$ 
3:  $K_R^{enc-new} \leftarrow \text{encrypt}_{K_{LI}^{new}}(K_R)$ 
4: update in  $F_{LI}$ :  $K_{LI}^{enc-new}, K_R^{enc-new}$ 
```

C. Combining Approaches

The approaches presented above can be composed, either sequentially or in parallel. By sequential composition, we mean that the user must *both* correctly answer security questions and receive e-mail. By parallel composition, we mean that either mechanism can be used alone to recover the password. The latter is achieved by using both systems in parallel.

For sequential composition, the user picks a uniformly random string r of the same length as K_{LI} . The user stores r in one of the mechanisms, and $K_{LI} \oplus r$ in the second mechanism, where \oplus denotes the exclusive-OR operation. If one recovers both of these, K_{LI} can be computed. If one learns only one of the pieces, nothing is gained, as both r and $K_{LI} \oplus r$ are uniformly random. More generally, to combine n mechanisms in arbitrary ways, (n, k) -secret sharing can be used. What we describe here are two trivial such schemes for $n = 2$.

VI. SECURITY

The goal we set is to emulate the security provided by a centralized solution. Some security risks are inherent to the password functionality, and apply regardless of implementation technique. For instance, in e-mail based recovery, an attacker compromising the victim's e-mail account can reset her password.

In this section, we elaborate on security concerns of our protocols. We do not have full cryptographic security proofs of our protocols, something which is important future work.

Concerning safety, we have designed our protocols such that persistently stored data remains in a consistent state if the protocol is aborted at any point. Some protocols may, if an operation fails, leave orphan files in the storage. If our protocol for revoking remembered credentials is interrupted, it may revoke more devices than intended. Apart from this, our operations have transactional semantics, assuming small writes (both creation, and updates) to the storage are atomic operations and that operations block until successful.

A. Adversary Model

To capture the concept of collusions, we consider an adversary that corrupts a number of nodes. Upon corruption, the adversary gains all information known to that node, and in the case of an active adversary, can also control its future actions. The adversary can also make requests to the underlying system, e.g., read files from the distributed storage. As almost all our protocols mainly operate on publicly readable (encrypted) data, this ability is important. The only computation made by nodes different from the one logging in are in verifying write operations, and in e-mail based password recovery.

B. Risks in Used Components

As our protocols make use of several standard components, vulnerabilities in those components can also affect our system. For instance, an adversary may prevent a user from logging in by attacking the victim's ability to read her login information file from the distributed storage. We note that there are many security techniques in DHTs that can mitigate such threats, and refer to Urdaneta et al. [24] for a recent survey.

C. Offline Guessing

An issue in password authentication based on a (distributed) file system is that the information required to verify a password attempt is inherently exposed. This means that in our scheme, we cannot prevent an attacker from mounting an offline attack against our encrypted passwords. This is a considerable drawback from centralized schemes, where the encrypted password database is kept protected.

As a partial mitigation, we utilize a KDF with a per-user salt. This forces an attacker to evaluate the KDF individually for each user on a password guess, defeating parallel attacks against multiple users. We recommend the system be instantiated with a slow KDF, such as bcrypt [25] to throttle offline guessing. The protocol could be modified to reduce storage by using the username as a salt, but we recommend against that as it would be vulnerable to pre-computation (before the system is started, or between instances using the same KDF) attacks against common username-password pairs.

The problem of a server performing offline attacks against its password database was treated by Ford and Kaliski [20]. Their techniques are client-server based, and require all servers to be online for a login. We leave it as future work to investigate modifying their protocol to be applicable also in a P2P setting. This would prevent offline guessing attacks.

D. Colluding Nodes in E-mail Based Recovery

In the mechanism for e-mail based recovery, we employ (n, k) -secret sharing, and secrets are stored on n random nodes in the system. If an attacker controls k or more of these nodes, she can recover the secret and access the victim's account. In Section VII-B we discuss the choice of these parameters.

A peer sampling protocol is used to select the n nodes where the shares are stored. An active attacker may influence this protocol, in order to ensure that she controls k out of the selected nodes. This can be mitigated by a peer sampling protocol designed to tolerate active attacks, such as Brahm's [16].

The protocol is designed to reveal only minimal information to the peers. Thus, even when colluding, peers have to get hold of both, the salt $ksalt_i$ and the recovery e-mail address *email* to mount an attack. The e-mail address will be revealed to a peer only when the user initiates the recovery process. Malicious peers might try to guess it earlier, but to verify guesses either $ksalt_i$ or $esalt_i$ are required. These are stored in the login information file of the user, which can only be found knowing the username. Therefore the setup process should be anonymous, where the peer does not learn the username of the user for whom it stores the share.

E. Updating Application Keys

When a password is changed, or device is revoked, access is effectively revoked from future updates to the key store. However, a malicious device may have stored the last keys it was able to access. Thus, in the password update procedure (which is also used for device revocation), the keys for the P2P application itself need to be updated, and then these new keys need to be written to the key store. How to update

them, if possible, is beyond the scope of this paper, as it is a functionality of the application protocol.

F. Denial-of-Service

An adversary may mount a DoS attack by writing many user names into the DHT, thus blocking those from registration by legitimate users. This attack is also possible in centralized systems, but there detection and counter-measures (e.g., removing the fake accounts) is significantly easier. One can solve this issue by assuming a lightweight CA dealing only with checking user identification before account creation, similar to Safebook [3]. We remark that this attack only target the availability of registration of a new user account, it does not affect existing users.

A related attack can occur when an adversary can prevent a user from accessing some of the data required to log in (e.g., by controlling all replicas holding the victim's entry in the DHT). In such attacks, existing users can be prevented from logging in, possibly permanently by overwriting or deleting the key. This illustrates the importance of applying security techniques for underlying components [24] and indicates a large replication factor should be chosen.

G. Security Summary

Aside from the concerns outlined, we believe that our schemes produce a similar level of security as client-server based password authentication schemes. The cryptographic design of our protocols relies on relatively standard techniques. This leads us to be confident that the security of our protocols can be formally proven using cryptographic techniques.

We also provide some features which are not commonly present in centralized systems. One of these is the ability to revoke stored credentials from only some specific devices. A second one is the ability to set up e-mail based password recovery without revealing your e-mail address before recovery actually occurs, offering an additional privacy protection.

VII. EVALUATION

We developed two lightweight custom simulators, one to evaluate the efficiency and security of our protocols, and one to assist in setting the n and k parameters for e-mail based password recovery. For the performance analysis, we take as input the time to perform required cryptographic operations, as well as the time of our network operations. For the analysis of n and k parameters, we need data on node uptime to evaluate the availability of the password recovery system for a given choice. We used latency data from Jiménez et al. [14], and node availability data from Rządca et al. [26].

To measure the computational cost of the necessary cryptographic operations in a prototype implementation, we used OpenSSL's built-in benchmark function on a 2.26 GHz Core 2 Duo running Mac OS X, as well as an ARM 1 GHz Cortex-A8, similar to modern smart phones. The times for all required cryptographic operations (using DSA as a public-key scheme) was negligible, below 5 ms (2 ms on the faster CPU).




Fig. 3. CDF for login latency in three modes: First time login, remembered logins, and first time login after password entry (pre-fetch). Network operations are assumed to have costs of BitTorrent mainline DHT lookups, using NR128-A (solid lines) or μ Torrent strategy (dashed lines) [14].

As our protocols can be applied with any DHT and file storage combination, we used the BitTorrent Mainline DHT as an example for our numeric performance evaluation. Jiménez et al. [14] recently ran experiments to evaluate the performance of their proposed algorithmic improvements. From their measurements, we received a CDF for the latency of real-world DHT lookups. We assume that all our *network operations*, writes and reads, both from DHT and distributed storage, take the same amount of time as a DHT lookup in their study. This can be motivated, as our distributed storage could be implemented via a DHT. We remark that their measurements are performed with a “warmed up” client with filled routing tables for the DHT. Thus, these numbers may be overly optimistic for a newly started client.

Finally, we believe, node availability will vary considerably between applications. As a representative case, we considered a distributed storage system by Rzacca et al. [26], which featured such data in their evaluation. In the distribution, 10% of nodes have availability 95%, 25% have availability 87%, 30% have availability 75%, and 35% have availability 33%. To this rough distribution, Gaussian noise with $\sigma = .1$ is added, and the resulting availability is capped between 3% and 97%.

A. Performance

We believe that the main performance-critical operation is logging in [27]. We believe that for all other operations, latency on the order of a few seconds can be acceptable, and even a minute if they are run in the background. Thus, we only present results for logging in, but note that as the operations for other protocols are similar, results are expected to be similar. The performance cost of our protocols is dominated by network operations. However, to slow down password guessing attempts, one may wish to force the key derivation procedure to be slow, to the point of making that cost dominant.

When evaluating the protocols, we parallelized network operations where possible. Logging in for the first time and remembering the credentials for future logins is then a sequence of two network operations, followed by key derivation,

TABLE III
LATENCIES OF PROTOCOLS, IN MILLISECONDS.

DHT	Network Op. [14]		First Login		Remem. Login	
	median	99 th	median	99 th	median	99 th
NR128-A	164	567	650	1362	346	915
μ Torrent	647	5140	3299	10154	1456	7148

followed by three parallel network operations. In a password login, it is also possible to pre-fetch some of the information after the user has entered her username, but before she enters her password. In particular, as soon as the username is known, the filename F_{LI} can be retrieved from the DHT, and the file can be read. Decryption of the file and further processing is then only possible after the user enters her password. To evaluate this speed-up, we computed the time it takes to finish the login after the user has entered her password. The time to fetch the two files is identical to the time to do a remembered login, and it is sufficiently small that the data can realistically be retrieved while the user is typing her password.

To determine the sensitivity of our performance to implementation characteristics, we evaluated our protocol for two different client strategies in the BitTorrent Mainline DHT: The NR128-A algorithm [14], and the μ Torrent client’s implementation. We present these performance numbers in Figure 3 and Table III. Firstly, we observe that with a fast storage, our login protocol is very fast, with a median login time of 650 ms the first time, and 346 ms for remembered logins. Comparing the results, we observe that our protocols are indeed sensitive to storage latency. While performance results building on μ Torrent data are slightly above recommended levels [28], we still consider them within range of acceptability for P2P.

When evaluating these numbers, we assumed that the runtime of the KDF function is negligible. As a system designer, one may wish to pick a slow KDF (e.g., bcrypt [25]), as this slows down password guessing attempts. Any latency intentionally added via the KDF would affect the first-time login (after the password entry) times.

B. Parameters for E-mail Password Recovery

There are trade-offs between availability, security, and storage space in our e-mail based password recovery protocol.

For the selection of k , the minimum number of peers required to recover the password, there is a direct trade-off between security and availability. Lower choices of k increase the risk of an adversary, controlling a significant number of nodes, to break into the user’s account. A higher k reduces the availability of the recovery functionality, which reflects the chances of a user to immediately succeed with the password recovery. However, if the user does not instantly receive k answers, she can simply wait until enough peers are online.

We believe a reasonable choice of parameters is $n = 16$ and $k = n/2$. With these numbers, using the availability data from Rzacca et al. [26], there is a 96% probability of immediate recovery of a lost password. A very strong attacker, corrupting 25% of the nodes in the system, would still only be able to

access the user's account with probability 3%. The analysis of parameter choice here is similar to any P2P system using secret sharing, and we refer to e.g., Vu et al. [11] for a more in-depth discussion.

C. Scalability

The latency of our protocols will scale similarly to DHTs or other distributed storage systems. The data we used for evaluation is based on measurements on the largest deployed DHT, demonstrating that performance is good with extremely large user numbers. Performance may in fact be worse for a small system, as there are then fewer nodes, meaning that it is less likely to find data at a nearby node. To bootstrap the system with good performance when it is small, a very simple distributed storage using one or a few super-nodes would be one approach. Storage requirements per user are also small, with a few files per user and small file sizes. From this, we conclude that our system is likely to scale well with the number of users.

VIII. CONCLUSIONS AND FUTURE WORK

The pros and cons of password-based authentication have been extensively debated. We believe that for some applications, a username-password pair provides an appropriate level of security. We argue that incorporating a well-known authentication scheme may assist in user adoption of P2P systems for more complex tasks than file sharing. To the best of our knowledge, ours is the first work to focus on password-based logins in a P2P setting, including mechanisms to recover a forgotten password. Our protocols are new (but our security questions are similar to [12]), relatively straightforward, and we believe, they are an important first step towards usable authentication in P2P.

The performance of our mechanisms in terms of delay varies according to the underlying DHT or P2P system in general and in relation to how much intentional delay is added by parameterizing cryptographic functions. Overall, however, our evaluation results show that for user satisfaction [27], the delays can be kept at a very acceptable level [28].

While we have provided an initial discussion of the security properties of our protocol here, future work will include a thorough security analysis. Our scheme allows offline password guessing attacks, which will also be addressed in future work.

ACKNOWLEDGMENTS

We thank Raúl Jiménez et al. [14] for sharing their measurement results and Jay Lorch for excellent work as a shepherd of this paper. This research has been funded by the Swedish Foundation for Strategic Research grant SSF FFL09-0086 and the Swedish Research Council grant VR 2009-3793.

REFERENCES

[1] C. Herley and P. C. van Oorschot, "A research agenda acknowledging the persistence of passwords," *IEEE Security & Privacy*, vol. 10, no. 1, pp. 28–36, 2012.

[2] T. Isdal, M. Piatek, A. Krishnamurthy, and T. E. Anderson, "Privacy-preserving P2P data sharing with OneSwarm," in *SIG COMM*. ACM, 2010, pp. 111–122.

[3] L. A. Cuttillo, R. Molva, and T. Strufe, "Safebook: A privacy-preserving online social network leveraging on real-life trust," *IEEE Communications Magazine*, vol. 47, no. 12, pp. 94–101, 2009.

[4] S. M. A. Abbas, J. A. Pouwelse, D. H. J. Epema, and H. J. Sips, "A gossip-based distributed social networking system," in *WETICE*, S. Reddy, Ed. IEEE Computer Society, 2009, pp. 93–98.

[5] K. Aberer, A. Datta, and M. Hauswirth, "Efficient, self-contained handling of identity in peer-to-peer systems," *IEEE Transactions on Knowledge and Data Engineering*, vol. 16, no. 7, pp. 858–869, 2004.

[6] J. R. Douceur, "The Sybil attack," in *IPTPS*, vol. 2429. Springer, 2002, pp. 251–260.

[7] M. Lillibridge, S. Elnikety, A. Birrell, M. Burrows, and M. Isard, "A cooperative internet backup scheme," in *USENIX*. USENIX, 2003, pp. 29–41.

[8] L. P. Cox, C. D. Murray, and B. D. Noble, "Pastiche: Making backup cheap and easy," in *OSDI*. USENIX Association, 2002.

[9] I. Clarke, O. Sandberg, M. Toseland, and V. Verendel, "Private communication through a network of trusted connections: The dark freenet," 2010. [Online]. Available: <https://freenetproject.org/papers/freenet-0.7.5-paper.pdf>

[10] K. Bennett, C. Grothoff, T. Horozov, and J. T. Lindgren, "An encoding for censorship-resistant sharing," 2003. [Online]. Available: <https://gnunet.org/svn/GNUnet-docs/WWW/download/ecrs.pdf>

[11] L.-H. Vu, K. Aberer, S. Buchegger, and A. Datta, "Enabling secure secret sharing in distributed online social networks," in *ACSAC*. IEEE Computer Society, 2009, pp. 419–428.

[12] N. Frykholm and A. Juels, "Error-tolerant password recovery," in *CCS*. ACM, 2001, pp. 1–9.

[13] K. Wehrle, S. Götz, and S. Rieche, "Distributed hash tables," in *Peer-to-Peer Systems and Applications*, vol. 3485. Springer, 2005, pp. 79–93.

[14] R. Jiménez, F. Osmani, and B. Knutsson, "Sub-second lookups on a large-scale Kademlia-based overlay," in *P2P*. IEEE Computer Society, 2011, pp. 82–91.

[15] M. Jelasity, S. Voulgaris, R. Guerraoui, A.-M. Kermarrec, and M. van Steen, "Gossip-based peer sampling," *ACM Computing Surveys*, vol. 25, no. 3, 2007.

[16] E. Bortnikov, M. Gurevich, I. Keidar, G. Kliot, and A. Shraer, "Brahms: Byzantine resilient random membership sampling," *Computer Networks*, vol. 53, no. 13, pp. 2340–2359, 2009.

[17] K. Bennett, C. Grothoff, T. Horozov, and I. Patrascu, "Efficient sharing of encrypted data," in *ACISP*, ser. Lecture Notes in Computer Science, L. M. Batten and J. Seberry, Eds., vol. 2384. Springer, 2002, pp. 107–120.

[18] S. C. Rhea, P. R. Eaton, D. Geels, H. Weatherspoon, B. Y. Zhao, and J. Kubiatowicz, "Pond: The OceanStore prototype," in *FAST*. USENIX, 2003.

[19] International Electrotechnical Commission, "ISO/IEC 27002:2005. Information technology – Security techniques – Code of practice for information security management," International Organization for Standardization, 2005.

[20] W. Ford and B. Kaliski Jr., "Server-assisted generation of a strong secret from a password," in *WET ICE*. IEEE Computer Society, 2000, pp. 176–180.

[21] J. G. Brainard, A. Juels, R. L. Rivest, M. Szydlo, and M. Yung, "Fourth-factor authentication: somebody you know," in *CCS*. ACM, 2006, pp. 168–178.

[22] G. Scoville, "Good security questions," <http://goodsecurityquestions.com/> (19th April, 2012).

[23] A. Rabkin, "Personal knowledge questions for fallback authentication: security questions in the era of Facebook," in *SOUPS*. ACM, 2008, pp. 13–23.

[24] G. Urdaneta, G. Pierre, and M. van Steen, "A survey of DHT security techniques," *ACM Computing Surveys*, vol. 43, no. 2, p. 8, 2011.

[25] N. Provos and D. Mazières, "A future-adaptable password scheme," in *USENIX Annual Technical Conference, FREENIX Track*. USENIX, 1999, pp. 81–91.

[26] K. Rzađca, A. Datta, and S. Buchegger, "Replica placement in P2P storage: Complexity and game theoretic analyses," in *ICDCS*. IEEE Computer Society, 2010, pp. 599–609.

[27] A. Rushinek and S. F. Rushinek, "What makes users happy?" *Commun. ACM*, vol. 29, no. 7, pp. 594–598, 1986.

[28] N. Tolia, D. G. Andersen, and M. Satyanarayanan, "Quantifying interactive user experience on thin clients," *IEEE Computer Society*, vol. 39, no. 3, pp. 46–52, 2006.

Summary Review Documentation for

“Passwords in Peer-to-Peer”

Authors: Gunnar Kreitz, Oleksandr Bodriagov, Benjamin Greschbach,
Guillermo Rodríguez-Cano, Sonja Buchegger

REVIEWER #1

I like this paper a lot, mostly because it tackles a useful problem and devises a highly practical solution. It is typical for papers about P2P systems to secure them by assuming each user possesses a cryptographic key, and this paper rightly points out that key management is not easily achievable for most users. By devising your solution as a way to derive cryptographic keys from passwords, you create a system that has general applicability.

I especially appreciated the explicit descriptions of each algorithm, that make the solution even clearer. These were particularly helpful in explaining the devmap, which could be explained a little more clearly in the text.

The main source of concern I have with the paper is that the solutions use fairly well-understood techniques in fairly obvious ways. Thus, I have the feeling that if a talented developer were given the task of taking a P2P system secured with keys and turning it into one secured by passwords, he might come up with these solutions without having to read the paper. In this respect, I see the main important contribution of your work as the statement of the problem. Nevertheless, it is useful that you codified the solution here, as this way anyone can have the benefit of your work.

Another issue I have is that there are no formal security proofs here, and it seems you don't have any prepared for inclusion in an accompanying tech report. It would be of great value to produce these, as it is typical for security systems to have subtle flaws that are only revealed when they are subjected to rigorous analysis.

I like that you evaluated performance, and got quite acceptable numbers. You did a good job of getting reasonable sources for distributions and performance numbers for components. However, the evaluation would be much better if you had actually built your system, so you could evaluate its performance directly. Building the system in earnest might also reveal interesting other properties of your algorithm, which you could report on for the benefit of the reader.

On a minor note, you could probably spend a lot less time in the evaluation section on graphing the effect of n and k on secret-sharing schemes. These types of analyses have appeared countless times before in papers on such schemes, and are well known. One or two text paragraphs, and no graphs, would probably be sufficient to talk about reasonable choices of n and k and their effect on probability of subversion.

It might be useful to remark about the fact that your e-mail recovery scheme requires peer software to be able to send e-mails. In today's era of spamming botnets, users may be appropriately reluctant to grant such capability to software they install. Still, it is probably reasonable to assume users

will allow this given the benefits it affords to the system at large.

Strengths: The paper tackles an important problem, securing P2P systems in a way that is cognizant of user limitations and desires. The paper is clearly presented, with algorithms not only described in text but laid out in explicit detail. The paper goes beyond merely enabling passwords to enabling several auxiliary requirements like password recovery.

Weaknesses: No proofs of security properties are given, and it does not seem the authors have them to put in an associated technical report. The solutions presented are fairly straightforward, not requiring any novel theoretical or systems research. So, they could probably be worked out by a practitioner without that practitioner having to read the paper. Too much space is spent in the evaluation section on well-known properties of the effectiveness of secret-sharing schemes.

REVIEWER #2

I like the paper's premise, but there are some significant issues that need to be addressed before the work is publishable. Especially the design issues (points (2)-(3) below) indicate to me that more work is required, even at the design level. There may be additional design issues that I have not spotted. Some mayor points: (1) The related work section is too short and leaves out important prior work. Even the original freenet implementation used a few similar ideas. (2) The salt in the `F_LI` message is redundant. For this reason, the salt should be dropped in order to simplify the system. (3) The password change mechanism seems to offer little protection against compromised passwords. It seems the password change mechanism really needs to require a change to the keys in the key-store in order to be effective.

Strengths: The authors' thinking is comprehensive and there is some merit in their vision.

Weaknesses: The work appears unfinished and contains several design flaws (some related to security), which I will outline below. The evaluation section focuses on performance, which seems less important. A better approach would be to focus on correctness and on the cost of attacking the system.

REVIEWER #3

This paper presents the p2p, decentralized emulation of centralized password login. The paper presents the scheme to support login based on users knowing a username-password pair.

Strengths: Clearly written and structured; an interesting topic and a straightforward, but neat idea.

Weaknesses: Why is it relevant in real world? why does this have to be p2p? How do you guarantee robustness and reliability? Compare your approach with p2p approaches using threshold cryptography for access control.

REVIEWER #4

Some positive points: (i) the paper explores a novel and creative idea, attempting to adapt the ordinary approach of password authentication to the context of P2P, (ii) it provides a detailed description of the solution, including algorithms, (iii) the performance evaluation shows that response times are acceptable (as long as DHT operation times are in fact as low as indicated), (iv) authors identify the main drawbacks of the solution and (v) the solution might be employed in future P2P applications for devices. Although the paper includes a set of algorithms and dedicates a lot of space to discuss implementation issues, such as required APIs, there is no implementation to actually demonstrate the feasibility of the solution. This would have allowed the authors to demonstrate both the feasibility of the assumptions and the actual performance in a small scale system. In terms of security, the solution allows the free access to the password database, which may be used in a brute force attack. The solution requires write operations on the DHT to be authenticated, but there is no proper explanation about how this is achieved, and implications on performance and vulnerabilities. Another requirement that affects the feasibility of the solution is the availability of an update operation: is this available on current DHT technology? If an update needs to be performed by remove + put, then the safety claims about the algorithms would not hold. A discussion on how trivial DoS attacks would impact on the scheme is missing (and how they could be avoided). Last, it is strange that there is no file delete operation in the distributed storage API, only file creation. How can a file be removed or the file system will grow forever?

Strengths: It seems the first attempt to bring conventional password authentication to P2P systems. Detailed description of the solution, including the required algorithms. Acceptable response times (as long as DHT operation times are in fact as low as indicated).

Weaknesses: Solution gives access to the password database, allowing brute force attacks to find out passwords (and authors state the attempt rate cannot be controlled). Seemingly strong requirements for the DHT such as update operations and authenticated write operations. Lacking an implementation and evaluation with a prototype.

REVIEWER #5

The authors assume some form of distributed storage outside the actual P2P system/DHT. I find that very artificial, why not then implement the whole password mechanism on it? And who maintains this storage? It is stated that the DHT has to be write-once on keys. But cannot individuals tamper with the code of the P2P client? On the adversary model, I would like to see a conclusion. The paper reads very well, and the bullet lists combined with the stars to show the difference

between the standard and the additional requirements are nice. However, some of the pseudocode (e.g., Algorithm 3) is very difficult to read. A good explanation/description often obviates the need for pseudocode. I would have liked to see something regarding replication. Will files be replicated by more than one peer? How will this be achieved? I think the claim that you cannot revoke the credentials of a single device is incorrect. Even from Gmail I can logout from another device/browser. If you mean excluding a device which later cannot have access anymore because the updated password is only known at the included devices, then indeed the solution might have some benefits.

Strengths: The paper presents a very detailed explanation of the protocols and steps involved for generating/login/changing etc. Most if not all "password" scenarios are included; creating, login, changing, logout, excluding a device, recovering using a security question, password recovery using an email address. Using figures from actual deployed systems, the login latency and security have been evaluated.

Weaknesses: I wonder whether we are solving a real problem here. Intellectually, it is interesting to have a password mechanism in P2P systems, but will we ever see P2P systems at a large scale that really use it? Section VI on security is weak. No actual statements on the security of the protocol are made, and in the conclusions (Section VIII), it is even stated that future work will include a thorough security analysis (see details below)!! This seems very strange in a paper on a password mechanism. The paper makes several significant but unrealistic assumptions, e.g., that a DHT can provide a method in which a username can only be registered once, and that a DHT can be made such that a user can always read his login. How is this done if there is churn and peers storing usernames are not available?

RESPONSE FROM THE AUTHORS

Many of the reviewers' comments and the excellent shepherding by Jay Lorch helped improve the paper. They gave us the possibility to extend the discussion of vulnerabilities, include more related work, and clarify important functionalities of the proposed system. We agree with the reviewers concerns regarding the formal security analysis and the prototype evaluation. We consider these to be important parts of our future work. Motivated by the reviewers concern about the vulnerabilities arising from reliance on DHTs, we extended the security discussion on these issues. We included a discussion of the risks in used components, such as denial-of-service attacks against the DHT and refer to related work with suggestions for mitigating these threats. Furthermore we discuss more explicitly the importance of updating application keys after a password change operation. Finally, we addressed several minor issues, such as explaining the purpose of the salt in F_{LI} , reducing the discussion in the evaluation section on the effectiveness of secret-sharing schemes and justifying the feasibility of DHT and storage assumptions.