

CS 255: Intro to Cryptography

*Prof. Dan Boneh*Due **Feb. 13, 11:59pm**

1 Introduction

In many software systems today, the primary weakness often lies in the user's password. This is especially apparent in light of recent security breaches that have highlighted some of the weak passwords people commonly use (e.g., 123456 or `password`). It is very important, then, that users choose strong passwords (or “passphrases”) to secure their accounts, but strong passwords can be long and unwieldy. Even more problematic, the user generally has many different services that use password authentication, and as a result, the user has to recall many different passwords.

One way for users to address this problem is to use a password manager, such as [LastPass](#) and [KeePass](#). Password managers make it very convenient for users to use a unique, strong password for each service that requires password authentication. However, given the sensitivity of the data contained in the password manager, it takes considerable care to store the information securely.

In this assignment, you will be writing a secure and efficient password manager. In your implementation, you will make use of various cryptographic primitives we have discussed in class—notably, authenticated encryption and collision-resistant hash functions. Because it is ill-advised to implement your own primitives in cryptography, you should use an established library: in this case, the [Stanford Javascript Crypto Library \(SJCL\)](#). We will provide starter code that contains a basic template, which you will be able to fill in to satisfy the functionality and security properties described below.

2 Secure Password Manager

2.1 Implementation details

In general, a password manager (also called a *keychain*) application will store its password database on disk, protected by a strong master password; in addition, while it is in use, it may store an “unlocked” representation of the database in memory, from which it can provide the password for each desired domain. Instead of implementing a full standalone password manager application, for this project you will only be responsible for the core library. Thus, you will not need to implement the interactive front-end for interacting with the password manager, nor will you need to actually write the contents to disk. Instead, you will simulate these functionalities by providing features to serialize and deserialize your data structures to string representations, so that it would be easy to complete a full password manager application by writing these representations to disk.

Your password manager will keep its in-memory password data in a key-value store (KVS), represented by a Javascript object whose keys correspond to domain names, and whose values correspond

to passwords for each domain. For example, a sample password manager instance might store the following information:

Key	Value
www.google.com	password
www.example.com	123456
www.amazon.com	6U)qA10By%3SZX\$o
www.ebay.com	guest

Naturally, writing this information to disk in the clear is not secure. In this assignment, you will need to preserve both the confidentiality and the integrity of the values in your KVS. Thus, you will be encrypting all the values (i.e, the passwords for different domains) using an authenticated encryption scheme, namely AES-GCM. In order to accommodate a potentially large number of entries in the password manager, you will encrypt and store each record individually in-memory. In other words, it is not appropriate to encrypt the entire KVS as a single blob. This way, you do not have to decrypt every entry in the password manager when fetching a single record.

We also do not want to leak any information about the domains the user has stored in the password manager. At the same time, we want to maintain the ability to search for the data corresponding to a specific domain. In this assignment, the KVS (Javascript object) storing the password data should have as its keys the HMAC of each domain name, rather than the actual domain name in the clear.¹ Then, to look up the data corresponding to domain name x , you first compute $\text{HMAC}(k, x)$, where k is the secret key for HMAC, and check whether the result exists as a key in the key-value store.

If you simply encrypt each domain/password pair in the KVS directly, your implementation will probably leak some information about the lengths of the passwords. We will not consider such implementations secure; rather, your implementation must prevent an adversary from learning any information about the password lengths. (To make this feasible, you may assume that the maximum length of any password is 64 bytes.)

You should use the password-based key derivation function, PBKDF2, to derive keys from the master password that the user specifies (see Section 2.3 for a more detailed discussion on PBKDF2). As we will see below, PBKDF2 is deliberately designed to be slow, and therefore you want to call it just once to derive a key. You can then derive the specific keys that you will actually use (for the HMAC for the domain names and the AES-GCM for the passwords) by evaluating HMAC at two different arbitrary values using the key derived from PBKDF2, and using the outputs as your two keys (this will *not* necessarily be secure for all MACs, but it does work with HMAC since it is also a secure PRF).

Note that a secure password manager is *not allowed* to include the master password (or even a hash of it or any other values that would leak information about it), or any of the secret keys used, in the serialized password database on disk. Additionally, you should assume that an adversary has access to your source code – this means you cannot embed any values you hope to hide from the adversary in your code.

Important note: In practice, there are many other considerations to take into account when designing an application like a secure password manager. For instance, as we will see below, we

¹Technically, you will need to use some serialized string representation (like hex or base64) of the HMAC value.

make no effort to thwart timing attacks or other side-channel attacks, and indeed this is extremely difficult using a language like Javascript that is so far removed from the underlying machine model. Along similar lines, it is always a good practice to erase secret keys in memory when the application is finished using them, in case the memory is subsequently exposed to the attacker. However, this is likewise infeasible in Javascript, as many standard features and library functions (including SJCL) leave data in memory that may leak secrets—both on the call stack and in garbage-collected structures on the heap. Thus, while this project provides a valuable proof-of-concept illustration of many important aspects of crypto implementation, it is very far from a complete picture, and you should *not* rely on the code you produce to be secure in practical settings.²

2.2 Threat model

When designing any system with security goals, it is important to specify a threat model: i.e., we must precisely define the power of the adversary, as well as the condition the adversary must satisfy in order to be considered to have “broken” the system. Thus, we will now specify the threat model for our secure password manager, in the form of a security game (of the same flavor as the PRF or CPA games). In particular, just as the CPA game allows the adversary to specify messages of its choice, our definition will seem to give the adversary a great deal of power over the contents of the password database. It is important to remember that we must make such strong assumptions when attempting to show that a system is secure for general use, because we have no idea under what circumstances it may end up being deployed.

Our security game proceeds as follows. As usual, the password manager will play the role of the challenger—interacting with another implicit party, the disk storage—while the adversary will make a series of adaptive queries that determine the behavior of the system. Some of the adversary’s queries may include a contingency for each of two “worlds”—as in worlds 0 and 1 in the CPA game—and, as in the CPA game, the “world bit” parameter, b , will determine which series of queries the challenger actually executes. Each query will take one of the following forms:

1. Specify values $\langle \text{domain}, \text{password}_0, \text{password}_1 \rangle$ to be added to the database. In world 0, the challenger must run the password manager algorithm to add the domain-password pair $\langle \text{domain}, \text{password}_0 \rangle$ to the database, while in world 1, the challenger must run the same algorithm, but on the pair $\langle \text{domain}, \text{password}_1 \rangle$.
2. Specify a key (domain) that the challenger must remove from the password database.
3. Specify that the challenger must serialize the database state to “disk”, whereupon the adversary will receive the entire result of the serialization, and will be able to replace it with an alternative of its choice (which must then immediately be deserialized by the challenger, running the password manager algorithm).
4. Specify a key (domain) for authentication, at which point the challenger (again running the password manager algorithms) must send the adversary the password corresponding to that domain.³

²For an additional perspective on the pitfalls of implementing crypto in languages like Javascript, see <http://www.matasano.com/articles/javascript-cryptography/>.

³This ability models the fact that in real life, the adversary may control the server running at the other endpoint of the client’s query, and therefore may see the result of a client password manager’s attempted authentication.

As in the PRF and CPA games, we say that the adversary wins the game if its probability of outputting 1 (for its guess of the world bit b) differs by a nonnegligible amount between world 0 and world 1. Unlike the PRF and CPA games, however,⁴ we will need an additional restriction for our security definition here. In particular, we will only allow adversaries whose queries are “admissible” in the following sense:

- Whenever the adversary makes a query (4), requiring the challenger to send its password for a given domain d , on its *last* query (1) setting the password for d (if there was such a query at all), it must have been the case that $\text{password}_0 = \text{password}_1$.

(Indeed, it is fairly easy to see that without this restriction, the adversary could trivially win the game—by query (1), it could cause the challenger to add a password for some domain d , under the adversary’s control, so that the password value differed between worlds 0 and 1; and then require, by query (4), that the challenger authenticate to it using domain d ’s password.)

Intuitively, this definition captures the fact that even if the adversary is able to exert substantial control over the contents of the password database—and even if it controls some malicious remote servers—it still cannot learn anything about the passwords in the database for any *other* servers.

For this project, you will not be required to give a formal proof that your system fulfills the strong security definition we have just stated. However, your system **should** be secure under this threat model (and a proof should exist, even though you do not have to produce it).

However, you should note that, to satisfy such a strong definition, there are a number of interesting attacks that you will have to defend against, most notably **swap attacks** and **rollback attacks**. In a swap attack, the adversary interchanges the values corresponding to different keys. For instance, the adversary might switch the entries for `www.google.com` and `www.evil.com`. Then, when the user (for whatever reason) tries to authenticate to `www.evil.com`, the user inadvertently provides its credentials for `www.google.com`. It should be easy to see that an adversary able to perform a swap attack can easily win the security game we outlined above. In your implementation, you must provide a defense against a swap attack.

In a rollback attack, the adversary can replace a record with a previous version of the record. For example, suppose the adversary was able to retrieve the KVS in the example above. At some later time, the user changes her/his password for `www.google.com` to `google_pwd`, which would update the value for `www.google.com` in the KVS. However, the adversary can replace this updated record with the previous record for `www.google.com`. Note that, as in the previous section, merely using authenticated encryption does not protect against this attack. Rather, in your implementation, you should compute a SHA-256 hash of the contents of the password manager. You can assume this hash value can be saved to a trusted storage medium (inaccessible to the adversary)—such as a flash drive on the user’s person. Whenever you load the password manager from disk, you should verify that the hash is valid. This way, you can be assured that the contents of the KVS have not been tampered with.

Depending on your design, your defense against rollback attacks might also turn out to protect against the swap attacks described earlier. However, you **must still implement** an explicit defense against swap attacks. In other words, the defenses you develop must work **independently** of one another. Even if a SHA-256 hash is not provided from trusted storage, your scheme must

⁴But similar to the CCA game.

be secure against an adversary that swaps two records.

2.3 Using PBKDF2

In the previous section, we did not give a precise formulation of the security properties we are assuming of PBKDF2. To give the full picture, we would need to work in a framework called the “random oracle model,” which would take us too far afield. Instead, this section will provide some practical guidelines on how to use PBKDF2.

Although we will discuss this in more detail towards the end of the course, it is important to remember that when deriving anything using passwords, we should always use a randomly generated salt. The recommended length of the salt is 64 bits, or even 128 bits if you want to be extra-safe. The key derived from PBKDF2 will be a function of the provided password and the salt. If you wish to derive the same key again in the future, such as when you’re loading the keychain from disk, you will need to store the salt in the clear.

The idea behind salting passwords is to prevent an *offline dictionary attack*, where an attacker derives keys for commonly used passwords, and tries to see if they work for different users’ keychains. Salting slows the attackers down since they will have to do a separate dictionary attack for each user. To make this even more difficult, PBKDF2 is deliberately designed to be a slow function (it should take about 0.5-1s per evaluation on your laptop) so as to rate limit an attacker trying to brute force it.⁵

3 API description

Here are descriptions of the functions you will need to implement. For each function, we also prescribe the run-time your solution must achieve (as a function of the number of entries n in the password database). We will assume that the input values (domain names and passwords) are of length $O(1)$, and regard each operation on an efficient dictionary/object/in-memory-key-value-store as a single step. Of course, if your solution is asymptotically more efficient than what we prescribe, that is acceptable.

3.1 `keychain.init(password)`

- **password:** password used to protect the keychain (**string**)
- No return value
- Run-time: $O(1)$

This method should create a new KVS. This function is also responsible for generating the necessary keys you need to provide for the various functionality of the password manager. Once initialized, the password manager should be in ready to support the other functionality described in the API.

⁵Of course, this only goes so far, and so if you use a password manager, you should make sure that you choose a long master password that is hard to brute force.

3.2 `keychain.load(password, representation, trustedDataCheck)`

- **password**: password used to authenticate keychain (**string**)
- **representation**: JSON encoded serialization of the keychain (**string**)
- **trustedDataCheck**: SHA-256 hash of the keychain; note that this is an **optional** parameter that is used to check integrity of the password manager (**string**)
- Returns: **boolean**
- Run-time: $O(n)$

This method loads the keychain state from a serialized representation. You can assume that **representation** is a valid serialization generated by a call to `keychain.dump()`. This function should verify that the given **password** is valid for the keychain. If the parameter **trustedDataCheck** is provided,⁶ this function should also affirm the integrity of the KVS. If tampering is detected, this function should throw an exception. If everything passes, the function should return **true** and the keychain object should be ready to support the other functionality described in the API. If the provided **password** is invalid, the function should return **false**.

If this method is called with the wrong master password, your code must return **false**, and no other queries `keychain.get`, `keychain.set`, or `keychain.remove` can be performed on the password manager unless the client calls `keychain.init` or successfully calls `keychain.load`. It is incorrect for your password manager to pretend like nothing is wrong when the wrong password is provided to `keychain.load`, and only later, fail to answer queries.

3.3 `keychain.dump()`

- Returns: an array consisting of two components, where the first is a JSON encoded serialization of the keychain and the second is a SHA-256 hash of the contents of the keychain (**array**)
- Run-time: $O(n)$

If the keychain has not been initialized or successfully loaded into memory, this method should return **null**. Otherwise, this method should create a JSON encoded serialization of the keychain, such that it may be loaded back into memory via a subsequent call to `keychain.load`. It should return an array consisting of this, and also of the SHA-256 hash of the keychain contents (which is to be stored in trusted storage, and used to prevent rollback attacks).

3.4 `keychain.set(name, value)`

- **name**: domain name of entry to add to the password manager (**string**)
- **value**: value associated with the given domain to store in the password manager (**string**)

⁶In Javascript, if an argument to a function is not provided, its value will be the special sentinel value **undefined**. You can test that a value is not **undefined** using the expression: `x !== undefined`.

- No return value
- Run-time: $O(1)$

If the keychain has not been initialized or successfully loaded into memory, this method should throw an exception: `throw "Keychain not initialized."` Otherwise, the method should insert the domain and associated data into the KVS. If the domain is already in the password manager, this method will update its value. Otherwise, it will create a new entry.

3.5 `keychain.get(name)`

- **name**: domain name of entry to fetch (**string**)
- **Return**: **string** (the value associated with the requested domain, **null** if not found)
- Run-time: $O(1)$

If the keychain has not been initialized or successfully loaded into memory, this method should throw an exception: `throw "Keychain not initialized."` If the requested domain is in the KVS, then this method should return the the saved data associated with the domain. If the requested domain is not in the KVS, then this method should return **null**.

3.6 `keychain.remove(name)`

- **name**: domain name of entry to fetch (**string**)
- **Return**: **boolean** (**true** if record with the specified name is found, **false** otherwise)
- Run-time: $O(1)$

If the keychain has not been initialized or successfully loaded into memory, this method should throw an exception: `throw "Keychain not initialized."` If the requested domain is in the KVS, then this method should remove the record from the KVS. The method returns **true** in this case. Otherwise, if the specified domain is not present, return **false**.

4 Hints and Summary

4.1 Setup Instructions

You should run your code using Node.js. The setup is fairly straightforward: if you don't have it already, install Node on your system by downloading it from <https://nodejs.org/en/download/>. After installing it, you should be able to run the `node` and `npm` commands in your command line. Then, extract the starter code and `cd` into the directory `proj1`. You will need to run `npm install` – this will install the dependencies specified in the `package.json` file locally, into a new directory named `node_modules` under `proj1`. You should now be all set. We have provided a simple test suite, which you can run using the command `npm test` from this directory.

Note that this test suite does not cover all of the properties we will test, and in particular, does not capture many of the security requirements. We will run a more exhaustive set of tests when grading.

4.2 Implementation Details

- All the code you will have to write will be in the file `password-manager.js`. Please do **not** write any code in another file, since our auto-grader will be assuming that you haven't.
- Your password manager will depend on the Stanford Javascript Crypto Library (SJCL) for its underlying crypto implementation. However, you **should not need** to call the SJCL functions directly (and our starter code does not include it directly). We have provided a support code library, `lib.js`, which provides wrappers for any SJCL functions that you should need, as well as some additional utility functions – mainly to deal with operations on bitarrays (which is the data structure that SJCL uses internally). Note that in your objects, you will not be able to use bitarrays directly as keys – you can use the provided utility functions to convert them to a Base64-encoded string representation and back.
- You can have a look at the tests being run in the file `test/test-password-manager.js`. You are always welcome to write more tests to make sure your implementation satisfies the requirements, but you are not required to, and we will not be grading your tests. The tests are written using the MochaJS framework (<https://mochajs.org/>) with Expect.js for assertions (<https://github.com/Automattic/expect.js>), and should be fairly readable.
- Serialization and deserialization of your password database (for the `load` and `dump`) functions should be done using JSON, via the standard Javascript APIs: `JSON.stringify` and `JSON.parse`. Further, the object that you serialize should have the key-value store that you use for your domain names and passwords stored internally in a field called `kvs` on the object (and you should use this field when you're restoring your keychain – in particular, if this field is modified, your object should have a modified set of domains and passwords). This is just to facilitate autograding. The last test in our suite is there to help you make sure that you're including this field (although it does *not* make sure you're using it when restoring).
- If your application detects tampering with any of its values at any point (like, say, a swap attack), it should throw an exception (thereby terminating the execution). We will not test what exception is thrown; it is fine to throw a string with an English description of the potential tampering.
- All functions should make *at most* one call to PBKDF2.
- The *only* thing you should assume about SHA-256 is that it is collision resistant.
- Although you are unlikely to need it, documentation for SJCL is provided here: <http://bitwiseshiftleft.github.io/sjcl/doc/>.

4.3 Summary of requirements

To summarize, you must implement a secure password manager that satisfies the following properties:

- The underlying in-memory data structure for the password manager should be a key-value store (Javascript object), where the keys correspond to domain names and the values correspond to the passwords for the given domain names.
- The password manager should be protected by a master password. Your implementation cannot include the master password (or any values that would leak information about it) in the serialized password database.
- When you need to derive a key from a password, you should use PBKDF2.
- Your system should satisfy the security properties described in the threat model (Section 2.2). In particular, you should defend against swap attacks and rollback attacks.
- You should implement all of the API functions (Section 3) with the parameters described there. Notably, your defenses for swap attacks and rollback attacks should be *independent*—your system must continue to be secure against swap attacks even if a hash value from trusted storage is not provided.
- You also need to submit answers to the short-answer questions from Section 5.

We will post an update on Piazza with instructions on how to submit your code and answers.

5 Short-answer questions

In addition to your implementation, please include answers to the following questions regarding your implementation. Your answers need not be long, but should include important details.

1. Briefly describe your method for preventing the adversary from learning information about the lengths of the passwords stored in your password manager.
2. Briefly describe your method for preventing swap attacks (Section 2.2). Provide an argument for why the attack is prevented in your scheme.
3. In our proposed defense against the rollback attack (Section 2.2), we assume that we can store the SHA-256 hash in a trusted location beyond the reach of an adversary. Is it necessary to assume that such a trusted location exists, in order to defend against rollback attacks? Briefly justify your answer.
4. Because HMAC is a deterministic MAC (it always produces the same output when run multiple times with the same input), we were able to look up a domain name using its HMAC value. There are also randomized MACs, that output a different tag every time, even when given the same input. Explain how you would do the look up if you had to use a randomized MAC instead of HMAC. Is there a performance penalty involved, and if so, what?
5. Our design leaks the number of records in the password manager. Describe an approach to reduce the information leaked about the number of records. Specifically, if there are k records, your scheme should only leak $\lfloor \log_2(k) \rfloor$ (that is, if k_1 and k_2 are such that $\lfloor \log_2(k_1) \rfloor = \lfloor \log_2(k_2) \rfloor$, the attacker should *not* be able to distinguish between a case where the true number of records is k_1 and another case where the true number of records is k_2).