




Zero Knowledge (About) Encryption: A Comparative Security Analysis of Three Cloud-based Password Managers

Matteo Scarlata 
ETH Zurich

Giovanni Torrisi [†]
USI, Lugano

Matilda Backendal [†]
USI, Lugano

Kenneth G. Paterson 
ETH Zurich

Abstract

Zero Knowledge Encryption is a term widely used by vendors of cloud-based password managers. Although it has no strict technical meaning, the term conveys the idea that the server, who stores encrypted password vaults on behalf of users, is unable to learn anything about the contents of those vaults. The security claims made by vendors imply that this should hold even if the server is fully malicious. This threat model is justified in practice by the high sensitivity of vault data, which makes password manager servers an attractive target for breaches (as evidenced by a history of attacks).

We examine the extent to which security against a fully malicious server holds true for three leading vendors who make the Zero Knowledge Encryption claim: Bitwarden, LastPass and Dashlane. Collectively, they have more than 60 million users and 23% market share. We present 12 distinct attacks against Bitwarden, 7 against LastPass and 6 against Dashlane. The attacks range in severity, from integrity violations of targeted user vaults to the complete compromise of all the vaults associated with an organisation. The majority of the attacks allow recovery of passwords. We have disclosed our findings to the vendors and remediation is underway.

Our attacks showcase the importance of considering the malicious server threat model for cloud-based password managers. Despite vendors' attempts to achieve security in this setting, we uncover several common design anti-patterns and cryptographic misconceptions that resulted in vulnerabilities. We discuss possible mitigations and also reflect more broadly on what can be learned from our analysis by developers of end-to-end encrypted systems.

1 Introduction

Despite the rise of alternative authentication methods, users today still have to deal with passwords, often numbering in the hundreds [45]. Password managers help to tame the problem

[†]Part of this work was done while Giovanni Torrisi and Matilda Backendal were at ETH Zurich.

by providing a tool to securely store passwords, reducing the challenge of remembering many passwords to remembering just the one “master password” for the password manager. Cloud-based password managers outsource the storage to a remote server under the control of a service provider. At an abstract level, a user's passwords are collected in a single object which is then encrypted by the user's client under a cryptographic key derived from the user's master password, creating an *encrypted vault*. The client then uploads the encrypted vault to the server. When a user wishes to access a password for a particular service, their client authenticates to the service, retrieves the encrypted vault, and decrypts it locally with a user-provided copy of the master password.

Importantly in solutions of this type, the service provider does not see the vault plaintext and therefore does not immediately learn the user's passwords or other sensitive data. This is akin to the situation with end-to-end encrypted (E2EE) cloud storage, and while the terms *E2EE* or *client-side encryption* are sometimes used by vendors in this space, the most commonly used term is *Zero Knowledge Encryption*. The term *Zero Knowledge* of course has a specific technical meaning in the context of interactive protocols, but here the term is being used with a different meaning, as we shall see.

The cloud-based approach has multiple advantages: users can access their encrypted vaults from multiple devices; vaults can store other sensitive information beyond passwords (e.g. credit card data, personal documents); and the service can be extended to allow sharing of sensitive data within a family, group or organisation. The “access from anywhere” feature creates work for vendors, who have to support access from web browsers as well as stand-alone applications running on different OSes. Many vendors have offerings which allow the cloud storage element to be self-hosted by an organisation instead of by the vendor.

Three prominent providers in this space are Bitwarden, Dashlane and LastPass. At the time of writing, Bitwarden claims to have 10 million users, Dashlane 19 million users and 24,000 business customers, and LastPass 33 million users and 100,000 business customers [17, 23, 42]. A 2024 report [49]

based on a survey of 1000 US consumers gives further insight into the popularity and market share of password managers. The built-in password managers of Google and Apple now represent 55% of the market, up from a combined share of only 15% in 2021. Bitwarden and LastPass are the next two largest, according to the study, with 11% and 10% market share, respectively. Dashlane now has only 2% market share, down from 7% in 2021 when it was amongst the market leaders.¹ There is a long tail of smaller players in the market.

Threat Model. We conduct our analysis of cloud-based password managers in a threat model in which the servers storing users' vaults are assumed to be fully malicious, meaning that they can arbitrarily deviate from expected behaviour when interacting with clients. This is a stronger model than considered in prior work. We present three arguments to justify this. First, vendors' explanations of Zero Knowledge Encryption and E2EE terminology imply that security is maintained in this setting. Second, given the large amount of sensitive data that they store, the service providers are likely targets for sophisticated attackers capable of penetrating the servers and then mounting active attacks (indeed, some vendors have been repeatedly breached [34, 39, 52]). Third, in closely related areas such as cloud storage and messaging, security against a malicious server is by now *de rigueur*.

Security Expectations. Given this threat model, what security should users reasonably expect of password managers? Based on prior work and our examination of vendors' security claims, we argue that the appropriate guarantees are:

1. Confidentiality of all data in a user's vault. This includes the *items* – that is, credentials (usernames and passwords), but also credit card information and notes. Most metadata is not kept confidential in the systems we studied, but achieving this is also desirable.
2. Integrity of vault data (and metadata). This should apply at the level of individual items, but also at the level of the data collection as a whole. It should be infeasible for an adversary to modify existing data or inject new data, undetectably delete or modify any data, or change the semantics of any data in a vault, e.g., by altering metadata. In particular, the integrity of any user-specific settings should be maintained.
3. Secure sharing and account recovery. If these features are enabled, then only the set of other users with whom a given user intends to share some or all of their vault data (or grant an account recovery capability) should be able to access that data (respectively, recover the account).

Several subsidiary security properties are needed to achieve these goals. For example, it should be hard for a malicious

¹ The apparent inconsistency between number of users and market share arises due to different sources: vendors' self-reported claims vs market study. Dashlane disputes the accuracy of this source due to the small sample size.

server to weaken the cryptography used by clients (e.g., reducing the iteration count used in password hashing).

More advanced goals can also be envisaged, e.g. vault indistinguishability, oblivious client operations, device privacy, etc. Such goals are far from being achieved by today's password managers and thus out of scope for our analysis; however, see [9] for a recent study of possible security properties.

Our Results. In Section 4, we give a detailed analysis of Bitwarden, Dashlane and LastPass, presenting a cornucopia of practical attacks. In the artefacts that accompany our paper, we give Proof of Concept (PoC) implementations of all of these attacks, demonstrating their feasibility. The attacks allow us to downgrade security guarantees, violate security expectations, and even fully compromise users' accounts. Table 1 lists the attacks and their impacts. Worryingly, the majority of the attacks allow recovery of passwords – the very thing that the password managers are meant to protect.

We group the attacks into four categories: attacks exploiting the key escrow features used for account recovery and SSO login, attacks based on lack of integrity of the vault as a whole, attacks enabled by the sharing features, and, finally, attacks exploiting backwards compatibility features.

These attacks reveal common design anti-patterns and cryptographic misconceptions. *Lack of authentication* of public keys is widespread. When combined with key escrow and sharing features, this results in the adversary being able to fully compromise vaults. Another recurring failure mode is (wrongly) assuming origin-authentication of public key ciphertexts, leading to *key substitution* attacks against Bitwarden in the style of [3, 36]. LastPass stands out for lacking any form of *ciphertext integrity*, using AES-CBC as its main encryption mode. Thanks to legacy code and backwards compatibility exploits, we can downgrade Bitwarden and Dashlane to similarly hazardous states. We also show that integrity is only achieved for single fields in individual items, instead of at the vault level. This enables cut and paste attacks within items and across the vault. Such attacks can often be chained to compromise the confidentiality of the vault as well. These attacks work even when proper authenticated encryption is used. They are possible because of insufficient key separation in vaults with complex structures and/or a lack of cryptographic binding between data and metadata.

Target Selection. We selected the three providers as targets based on their historical and current market share, the availability of unobfuscated client source code (available for Bitwarden and, partially, for Dashlane and LastPass), the richness of the offered feature set, and the diversity of approaches. As noted above, Apple and Google now dominate the password manager market. However, the unavailability of source code precludes easy analysis of these systems. Additionally, E2EE is an opt-in feature for Google's password manager, so users do not enjoy security against malicious servers by default. Finally, while numerous other cloud-based password

Class	Ref	Name	Cause	Impact	Interaction	Mitigations
Key Escrow	BW01	Malicious Auto-Enrolment	Lack of Key Auth, Key Substitution	Full vault compromise †	1 join	PKA, SC
	BW02	Malicious Key Rotation	Key Substitution	Full vault compromise †	1 rotation	SC
	BW03	Malicious KC Conversion	Lack of Key Auth, Key Substitution	Full vault compromise †	1 dialog	PKA, SC
Item-Level Vault Encryption	LP01	Malicious Password Reset	Lack of Key Auth	Full vault compromise †	1 login	PKA, SC
	BW04	Unprotected Item Metadata	Lack of Auth Enc	Read and modify metadata	-	AE, AD
	BW05	Item Field Swapping	Lack of Key Sep	Field and item swapping	-	KS
	BW06	Icon URL Item Decryption	Lack of Key Sep	Loss of confidentiality †	1 open	KS
	BW07	Remove KDF Iterations	Lack of Auth Enc	No brute-force protection	1 login	AD
	LP02	Item Field Swapping	Lack of Auth Enc	Field and item swapping	-	KS, AE
	LP03	Icon URL Item Decryption	Lack of Key Sep	Loss of confidentiality †	1 open	KS, AE
	LP04	Remove KDF Iterations	Lack of Auth Enc	No brute-force protection	1 login	AD
	LP05	Malleable Vault	Lack of Auth Enc	Loss of vault integrity	-	AE
	LP06	Unprotected Item Metadata	Lack of Auth Enc	Read and modify metadata	-	AD, AE
	DL01	Transaction Replay	Lack of Key Sep	Loss of vault integrity	-	KS
Sharing	BW08	Organisation Injection	Lack of Key Auth	Add users to arbitrary orgs	1 sync	SC
	BW09	Organisation Overwrite	Lack of Key Auth, Key Substitution	Organisation compromise †	1 join	PKA, SC
	LP07	Sharing Key Overwrite	Lack of Key Auth	Shared vault compromise †	1 join	PKA
Backwards Compatibility	DL02	Sharing Key Overwrite	Lack of Key Auth	Shared vault compromise †	1 join	PKA
	BW10	Disable Per-Item Keys	Lack of Auth Enc	Downgrade key hierarchy	-	KS, AD
	BW11	User Key Overwriting	CBC Support	Loss of confidentiality †	2 logins	AE
	BW12	Downgrade to Legacy	CBC Support	Full vault compromise †	2 logins	AE
	DL03	Item Injection	CBC Support	Loss of vault integrity	10 ⁴ syncs	AE
	DL04	Remove KDF Iterations	CBC Support	No brute-force protection	10 ⁴ syncs	AE
	DL05	CBC-Only Downgrade	CBC Support	Loss of confidentiality †	10 ⁵ syncs	AE
DL06	Lucky 64	CBC Support	No brute-force protection	10 ⁴ syncs	AE	

Table 1: Summary of attacks. The attacks are grouped in four categories, depending on the password manager feature they exploit. The attack ref(erence) also indicates the product: BW for [Bitwarden](#), LP for [LastPass](#), and DL for [Dashlane](#). We highlight the main causes for the attack, its high-level impact († denotes recovery of encrypted passwords) and the client interaction required: periodic or user-triggered synchronisation (sync); the user logging in (login), opening the vault (open), joining an organisation (join), sharing a vault (share) or clicking on a misleading dialog (dialog). Finally, we report the mitigations from Section 5.

managers could also have been studied, our deep examination of the three selected products – representing over 60 million users and 23% of the market – was already sufficient to surface severe design flaws and common misunderstandings.

Mitigations and Disclosure. We describe potential mitigations for each attack in Section 5. We also discuss general strategies to avoid the common pitfalls we discovered. We have shared these mitigations as part of a detailed, vendor-specific disclosure document with Bitwarden, LastPass, and Dashlane. In each case we have followed an industry-standard coordinated vulnerability disclosure approach, giving vendors at least 90 days to deploy countermeasures. For more details, see our ethics analysis in Appendix A.

Broader Impact and Contributions. The immediate impact of our work is to help the affected vendors patch their vulnerabilities, thereby providing stronger security for very

sensitive data of millions of users. But the impact of our findings also goes beyond the specific products they affect to illustrate some of the frontiers in the space of end-to-end encrypted applications. The importance of the integrity of vaults is a new finding, but parallels can be drawn with similar, recent notions of integrity for cloud storage [10]. Building systems that support key escrow without opening up attack avenues for malicious actors is considered a hard problem, one that has arisen repeatedly in the arena of lawful access and the Crypto Wars [1]. The problem of reliably authenticating public keys that (as we show) plagues password managers also comes up in the messaging space, where it is being tackled by the introduction of Auditable Key Directories, cf. [43]. Section 6 considers these and other issues at greater length.

Beyond categorizing the common pitfalls of E2EE cloud-based password managers and drawing parallels to where these arise in other E2EE systems, a main contribution of

our work is to firmly establish the need for formal study of these products and their security. The number of severe vulnerabilities we uncover among a diverse set of vendors shows that getting end-to-end encryption right in this setting is not easy. Building on initial notions for password manager security in [33], we propose as future work to formally define the goals of E2EE against a malicious server, akin to what was recently done for cloud storage [10]. Analyzing cloud-based password managers in this threat model is a novel aspect of our work, and we have seen from our interactions with vendors that our attacks provide a strong motivation to work towards achieving theoretically well-founded guarantees in this setting, promoting the benefits of formal analysis and provable security in practice.

Related Work. There was substantial prior analysis of password managers shortly after they first appeared on the market [12, 31, 33, 44, 50, 51, 55]. However, none of this work considered the malicious server threat model as we do. In 2020, [20, 46] revisited this early work concluding that password managers have improved their security. More recently, [30] developed injection attacks, their threat model assuming a malicious client in combination with a passive server (or a network adversary); [28] studied 43 password managers and their vulnerability to offline guessing attacks.

A parallel preprint [9] studies password managers in a threat model where the server remains “covert” in any attack it mounts. The paper classified cloud-based password managers into three types and sketched attacks by type. The attacks violate vault integrity, as well as several weaker properties not currently targeted by vendors. The attacks’ impacts are not fully explored, and there are no PoCs. Some of them rely on weak passwords and/or low PBKDF iteration counts.

We discuss related work in more detail in the full version of this paper [48].

2 Threat Model

As we laid out in our discussion of related work, there has been a significant amount of security analysis of password managers over the last decade (and more). At the same time, the landscape of password managers has changed, with cloud-based solutions becoming significantly more common. This makes it timely to revisit the question: what is the appropriate threat model in this setting?

Before going further, we note that there is a trivial attack against end-to-end encrypted applications that we wish to take out of consideration from the outset. Namely, a user can be supplied with a malicious client which, for example, simply returns the user’s master password to the server in the clear. Such an attack would be detectable through code audits and – though possible via an obfuscated functionality – would be very risky for a vendor to conduct. This is especially true for signed software releases (e.g., desktop and mobile clients,

browser extensions), as the existence of a signed malicious client would constitute proof of vendor misconduct.

Setting this trivial attack aside, we argue here that it is appropriate to consider a threat model in which the server is considered fully malicious, meaning that it can deviate arbitrarily from its expected behaviour. Such deviations are hard to detect, especially when targeted at selected users. We argue that this threat model is motivated by the advertised security guarantees by the vendors, the sensitivity of vault data, and the fact that malicious-server security is by now the norm for outsourced data in other, related settings. We provide example statements made by the providers which we believe justify this threat model here (taken from their web sites and white papers where they explain their “Zero Knowledge Encryption” claims), and discuss further arguments in Appendix C.

“Zero knowledge encryption: Bitwarden team members cannot see your passwords. Your data remains end-to-end encrypted with your individual email and master password. [...] Since it’s fully encrypted before it ever leaves your device, only you have access to your data. Not even the team at Bitwarden can read your data (even if we wanted to).” [18]

“Dashlane Password Manager is designed using zero-knowledge architecture, with the data encrypted locally on the user’s device. Only the user can access the data by using a password or another form of authentication. Since Dashlane doesn’t have access to the user’s vault and doesn’t store the user’s Master Password, malicious actors can’t steal the information, even if Dashlane’s servers are compromised.” [24]

“With a zero-knowledge approach, you can rest easy knowing that no one else but you, not even your password manager vendor, has the keys to the kingdom. [...] For example, zero-knowledge means that no one has access to your master password for LastPass or the data stored in your LastPass vault, except you (not even LastPass).” [26]

While none of the vendors quoted here explicitly describe the threat model under which they expect their product to remain secure in terms of attacker capabilities, we argue that – based on statements like the ones above – their customers could reasonably deduce that their vault data remains protected even against a malicious server. In addition to these general statements, the vendors to whom we have disclosed so far (Bitwarden, LastPass and Dashlane) have not questioned our adoption of this threat model for our analysis.²

Client Interaction. While our threat model concerns a malicious server, some of our attacks rely on interaction with the (honest) client victim. We treat the amount of client interaction needed for an attack to succeed as a resource, and for each attack, we list the expected client interactions required.

The majority of our attacks require simple interactions which users or their clients perform routinely as part of their

² LastPass have communicated that they don’t adopt a generic malicious server threat model, but apply compensating controls to mitigate server risks.

usage of the product, such as logging in to their account, opening the vault and viewing the items, or performing periodic synchronization of data. We also present attacks that require more complex user actions, such as key rotations, joining an organization, sharing credentials, or even clicking on a misleading dialog. Although assessing the probability of these actions is challenging, we believe that, within a vast user base, many users will likely perform them. Table 1 summarises the interaction requirements for each attack.

3 The Password Managers

3.1 Bitwarden

The key hierarchy of a Bitwarden user, illustrated in Figure 1, is rooted in the user’s master password and email. These are used respectively as key and salt of a KDF, which can either be PBKDF2 or Argon2id, in order to generate a 32-byte-long *master key* (k_m). The master key is used for authentication by sending the result (h_m) of a single iteration of PBKDF2 on the master key with the user’s password as salt to the server, which compares it to the value of h_m stored at registration. The master key is additionally used to derive two 32-byte keys via HKDF-Expand with two different input labels: “enc” and “mac”. The resulting keys (k_{em}^{enc} , k_{em}^{mac}) are concatenated to form the 64-byte *extended master key* (k_{em}).

After deriving the extended master key, a 64-byte *user key* (k_u) is sampled uniformly at random. Similarly to k_{em} , this key can be interpreted as consisting of an encryption (k_u^{enc}) and a MAC key (k_u^{mac}), each 32-bytes. The role of the user key is crucial: it protects all the data inside the user’s vault, either by directly encrypting vault items, or indirectly (if so-called “per-item keys” are enabled, as explained below). Once the user key is generated, it is encrypted under k_{em} and stored on the server to be served at the next login.

Additionally, the client generates an RSA keypair (with a 2048-bit-long modulus) consisting of a public key pk and a private key sk . This keypair is used in the context of Bitwarden’s organisation as described below. Finally, the private key is also encrypted under the user key and sent to the server.

Bitwarden allows users and businesses to self-host the server, in order to have complete control of their data. For users of self-hosted Bitwarden instances, the adversary in our threat model will thus be the self-hosted server.

Key and item encryption Bitwarden uses AES-CBC-HMAC for authenticated encryption of keys and vault items, using the “enc” and “mac” parts of the key for encryption and MACing, respectively. For backwards compatibility, the HMAC is omitted if the “mac” key part is missing. That is, if the client attempts to encrypt using a key which is only 32 bytes, the client falls back to using AES-CBC.

In vault items, each item data field is encrypted individually. Item fields used to be encrypted directly with the user key k_u ,

but from version 2024.2.0 new items are encrypted with individually sampled *per-item* keys, which are in turn encrypted by the user key. Old items remain directly under k_u .

Organisations Credential sharing is implemented in Bitwarden through organisations. When a user creates an organisation, their client samples a new 64-byte *organisation symmetric key* k_{org} , and encrypts this key under the user’s public key, yielding c_{org}^{user} . Additionally, the client also generates an *organisation RSA keypair* (pk_{org}, sk_{org}) (used for account recovery purposes), and encrypts the private key using k_{org} to create c_{skOrg} (see Figure 1). It then sends c_{org}^{user} , pk_{org} , and the encrypted c_{skOrg} to the server. Items in an organisation are individually encrypted using the organisation symmetric key.

The creator of an organisation can add new members by encrypting the organisation symmetric key for the public key of the invitee. For each user $user'$, the resulting ciphertext $c_{org}^{user'}$ is sent to the server for storage, and served to $user'$ when syncing. The client of $user'$ can then decrypt $c_{org}^{user'}$ using their private key to obtain k_{org} . A client determines that the user is a member of an organisation if it can obtain k_{org} .

Further organisation data, such as the organisation name and policies regarding account recovery, for example, are stored unencrypted on the server and sent to the client whenever a sync operation is performed.

3.2 LastPass

Similarly to in Bitwarden, the key hierarchy in LastPass (illustrated in Figure 1) begins with the user’s password, which is used to derive a 32-byte *user key* k_u using PBKDF2 with the user’s email address as salt. The default KDF iteration count is 600,000. To authenticate, the client sends h_u , the result of one more iteration of PBKDF2 over k_u with the password as salt, to the server; the server checks if the received h_u matches the value received at registration.

The user key is used to directly encrypt the items in the user’s vault using AES-CBC (with no integrity protection). Each user also has an RSA key pair (sk, pk); the secret key is encrypted under k_u using AES-CBC and stored at the server.

Sharing among LastPass users is implemented through shared folders. When a user first shares an item, their client samples a 32-byte shared folder key k_{shr} . This is subsequently used to encrypt the items in the shared folder using AES-CBC. The folder key is itself encrypted using RSA-OAEP under the public key of every user with whom the folder is shared.

3.3 Dashlane

Dashlane’s vault is maintained in the form of a transactional database. Each transaction can create, modify, or delete an item in the database: the transaction type and timestamp are unencrypted metadata fields, while the transaction content,

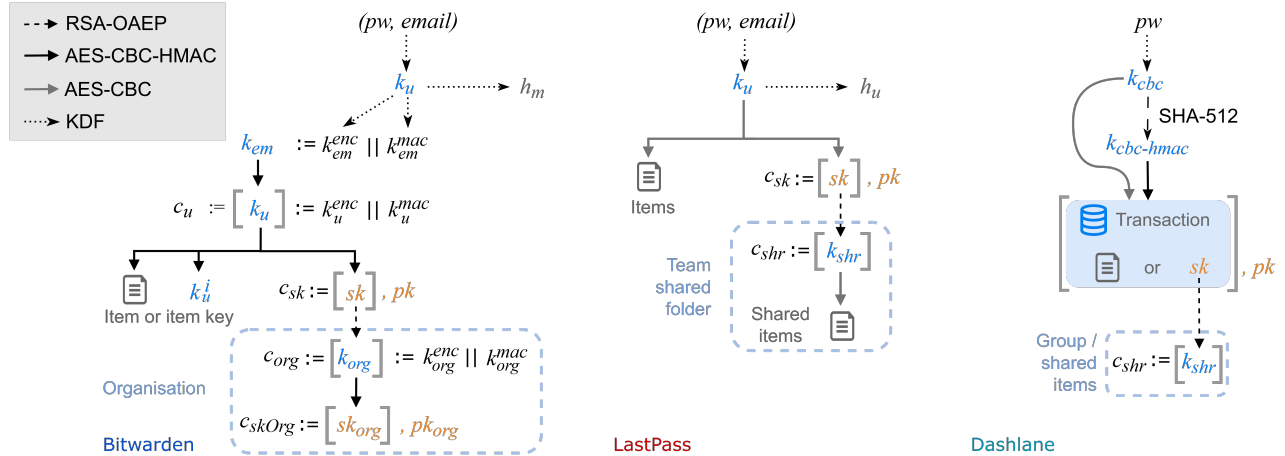


Figure 1: Key hierarchies for Bitwarden (left), LastPass (center) and Dashlane (right).

consisting of compressed XML data with information about items or user settings, is encrypted.

Content is encrypted using a key derived from the user’s master password. The exact details for these operations are specified in the first bytes of the content by a *payload* string: this defines the password-based key derivation function (either PBKDF2, Argon2d, or noderivation), its parameters (e.g., salt length, hash method, and iterations for PBKDF2), the cipher (by default AES-256), and the cipher mode (by default CBC with HMAC). First, an intermediate 32-byte key k_{cbc} is derived from the password. If the cipher mode is CBC with HMAC, then an additional SHA512 is used to derive the 64-byte $k_{cbc-hmac}$. When the user modifies their vault (e.g., by adding an item), a new transaction is created and sent to the server. See the illustration in Figure 1.

Each user also has a 2048-bit modulus RSA key pair (sk, pk) ; the secret key is stored encrypted (as a transaction) on the server, together with the public key. The RSA keys are used for sharing: when a user shares an item with another user (either directly or in a group), a symmetric encryption key k_{shr} is sampled and encrypted for the recipient(s) public keys (which the sharer retrieves from the server).

NOTE. The only fully-available open-source client is the CLI client, which does not support encryption nor the credential sharing feature, so the description we provide here is based in part on reverse engineering and may be incomplete. We focus on the parts of the system relevant to our attacks, for which the functionality we describe here is verified by the PoCs.

4 Attacks

We now present our attacks on Bitwarden, LastPass and Dashlane, grouped into four categories based on cause, as follows:

1. In E2EE password managers, forgetting the master password means that a user irrecoverably loses access to their

vault. Thus, many password managers implement some form of account recovery mechanism through **key escrow**, such that an administrator can help users recover their vault contents. In Section 4.1, we show attacks against vault confidentiality in Bitwarden and LastPass resulting from vulnerabilities in their key escrow designs. We also discuss a weakness in Dashlane’s design.

2. None of the password managers we analyse treat the password vault as a monolithic object. Rather, all data items (as well as sensitive user settings) are encrypted as separate objects, and often combined with unencrypted or unauthenticated metadata. In Section 4.2, we show how such an **item-level vault encryption** design is often in tension with the *vault integrity* security goal. We find attacks against all of Bitwarden, LastPass and Dashlane. Notably, many of these integrity violations also lead to concrete attacks impacting confidentiality.
3. Many password managers implement **sharing** features (targeting for example families and businesses) which allow multiple people to access common shared credentials. In Section 4.3, we show attacks against the sharing features of all three of our targets. These impact both vault integrity and confidentiality.
4. All three password managers that we analyse have undergone cryptographic evolution over time. However, they all maintain some form of **backwards compatibility** to support older client versions. In Section 4.4, we show that such “legacy code” creates cryptographic hazards. Specifically, a malicious server can trigger legacy code paths to be executed in clients, resulting in downgrade attacks in Bitwarden and Dashlane. This holds even if a user intends to use the latest version of the clients. Through such attacks, an adversary can for example downgrade clients to use unauthenticated encryption,

opening up avenues for further attacks that violate both vault integrity and confidentiality.

Table 1 gives a complete summary of our attacks. Throughout this section, we refer to “the adversary” as an entity that has full control of the server. We summarise the analysed software versions and the availability of source code in Table 2 in the full version of this paper [48]. Unless otherwise specified, all of the attacks affect all the clients mentioned there.³

Analysis Method: Our approach to analysing password managers follows the “attack” branch of the “Cryptography in the Wild” methodology that is described in detail in [5]. We explain here how we deployed this methodology in the specific context of password managers.

Our process began with reviewing white-papers and client code (and server code, when available). We initially focused on Bitwarden, where the availability of source code eased our analysis and enabled us to more easily build our mental model of architecture, protocols, and system features (cf. *object selection* in [5]). From the outset, our goal was to carry out an analysis in the face of a malicious or compromised server: based on the functionality of the target systems, we refined this to a concrete *adversarial model*. We developed pseudo-code models for the main cryptographic actions of each feature of interest (*ingestion*). We then analysed these models to find *attacks*. Here we relied on our own experience in finding cryptographic vulnerabilities in systems, along with a list of potential attack vectors gained from the by-now-extensive literature on cryptographic attacks. For password managers, we drew particular inspiration from recent literature on E2EE cloud storage providers in the malicious server setting [3, 4, 11, 35, 37] as well as common attack vectors such as padding oracle attacks [54], key overwriting in various forms, cf. [19], and exploitation of backwards compatibility features, cf. [38]. Finally, we built *PoC* implementations for the attacks, implementing malicious server functionality where needed, and testing against the relevant client. This step is vital for ensuring that attacks developed “on paper” actually work in practice.

Following our initial study of Bitwarden, we iterated upon the same process to determine whether our findings were applicable to other products. For Dashlane and LastPass, we carried out some reverse engineering of minified client Javascript and of the Android applications. Having already studied Bitwarden helped us in developing our understanding how these products operate. We confirmed that our attacks were applicable to the newly selected targets. Additionally, we manually classified the different attacks and grouped them into four categories, providing a starting point for follow-up work analysing further password manager products. This

³ Due to the considerable engineering effort required for testing all the clients, we limit our PoC testing to Bitwarden’s web client and the browser extensions for LastPass and Dashlane. For most of the attacks, it is evident that they are design issues, independent of specific client implementations.

extends the methodology from [5].

Note that absent formal analyses, the set of attacks we found must be considered incomplete. Indeed, in some cases, we stopped looking for attacks after accumulating sufficiently many severe enough vulnerabilities to demonstrate that the given target was vulnerable in our threat model. Future work may involve employing formal tools to conclusively show other vulnerabilities, or the lack thereof.

4.1 Key Escrow Attacks

These attacks all concern various forms of key escrow functionalities provided by the password managers to enable account recovery in the case of master password loss. Successful attacks on these recovery mechanisms have severe consequences. In most cases, the adversary can recover the full vault of the user, therefore compromising all of the data that it contains. Some recovery mechanisms are less powerful than others, and rely on storing recovery data on the user’s devices. In these cases, our attacks are also less impactful.

4.1.1 Bitwarden

Bitwarden supports two flavours of key escrow. The first is the *organisation account recovery* feature.⁴ This allows the admins of organisations to reset the master password of all members who enrol in recovery. Enrolling in recovery can either be a manual step requiring user interaction (so-called *self-enrolment*), or happen automatically when the user joins the organisation (*auto-enrolment*). When a user enrolls (by either method) in account recovery, the client fetches the organisation public key pk_{org} from the server, encrypts the user key k_u under pk_{org} , and sends the resulting *account recovery ciphertext* c_{rec} to the server. We describe two key recovery attacks caused by the account recovery feature here. The second flavour of key escrow in Bitwarden (the so-called “Key Connector” feature) is also vulnerable to a master key recovery attack (**BW03**). For space reasons, the description of this feature and the attack are deferred to the full version [48].

BW01: Malicious Auto-Enrolment. Recall that when a Bitwarden user joins an organisation, the client fetches the organisation’s data, such as name, key material, and organisational policies, from the server. This data is not integrity protected. As a consequence, an adversary controlling the server can compromise users as soon as they accept an invitation to join *any* organisation, as follows.

When a user accepts an invitation, the client asks the server for the account recovery policy and the public-key of the organisation. The adversary replaces the organisation’s real data, setting auto-enrolment to true in the policy, and replacing the public key pk_{org} with a malicious pk_{org}^{adv} for which

⁴ While the recovery feature is restricted to “enterprise” organisations (with a paid subscription), all Bitwarden clients include this feature’s code, and our attacks succeed independently of user subscription status.

they know the secret key sk_{org}^{adv} . Since account recovery is enabled, the client encrypts the user key k_u under the organisation public key pk_{org}^{adv} , and sends the resulting account recovery ciphertext c_{rec} to the server. The adversary decrypts c_{rec} with sk_{org}^{adv} and recovers k_u .

IMPACT. With possession of the user key, the adversary is able to read and modify all vault data of the targeted user.

REQUIREMENTS. The user joins any organisation. Note that the adversary does **not** have to control the organisation; it can be an honest invitation from a trusted enterprise or family member, for example.

NOTE. An organisation’s private key sk_{org} is known to all its members. It follows that if the adversary compromises a user, they learn the sk_{org} for all the organisations that user is part of. If any of these organisations enabled account recovery, the adversary then also compromises all the users of these organisations. This process can be repeated in a *worm-like* fashion, infecting all organisations that have key recovery enabled and have overlapping members.

BW02: Malicious Key Rotation. Bitwarden allows its users to change their master password, and optionally rotate their keys. If a user decides to rotate their keys, the client samples a new user key k'_u and encrypts it under the new extended master key k'_{em} (cf. Figure 1).

If a user is enrolled in account recovery when they rotate their keys, the client will also have to regenerate their account recovery ciphertext. To do so, it encrypts the new user key k'_u under the organisation public key pk_{org} , and sends the new account recovery ciphertext c'_{rec} to the server.

The key point here is that pk_{org} is not retrieved from the user’s vault; rather, the client performs a sync operation with the server to obtain it. Crucially, the organisation data provided by this sync operation is not authenticated in any way (see Attack [BW01](#)). This thus provides the adversary with another opportunity to obtain a victim’s user key, by supplying a new public key pk_{org}^{adv} for which they know sk_{org}^{adv} and setting the account recovery enrolment to true. The client will then send an account recovery ciphertext c'_{rec} containing the new user key, which the adversary can decrypt to obtain k'_u .

Note that this attack works whether or not the user is a member of an organisation at the time of key rotation: the adversary can either use Attack [BW08](#) or [BW09](#) to, respectively, add the user to a newly created organisation or overwrite an existing organisation’s keys. We describe the details of the attack in the full version [48].

IMPACT. With possession of the user key, the adversary is able to read and modify all vault data of the target user.

REQUIREMENTS. The user rotates their encryption keys.

4.1.2 LastPass

LastPass also provides two forms of key escrow: account recovery and admin *password reset*. We focus here on the

second. The master password reset is a feature for LastPass’ enterprise products “business” and “teams”.⁵ Within a team, privileged user accounts called *super admins* have the ability to reset the master password of users in the team. To this end, each admin account is associated with an RSA keypair (sk_{adm}, pk_{adm}) . After a super admin chooses to reset a user’s password, the client of the selected user will – upon the next login in the LastPass browser extension – retrieve the list of admins and their public keys from the server, encrypt its user key k_u under each admin’s public key pk_{adm}^i , and send the resulting ciphertexts back to the server.⁶

LP01: Malicious Password Reset. Master password reset relies on admin public keys pk_{adm}^i being retrieved from the server. The LastPass client does not authenticate these keys in any way. It follows that an adversary can trivially inject their own pk_{adm}^{adv} , for which they know sk_{adm}^{adv} . They can then recover k_u from the ciphertext sent by the client.

In theory, only users in teams where password reset is enabled and who are selected for reset should be affected by this vulnerability. In practice, however, LastPass clients query the server at each login and fetch a list of admin keys. They then send the account recovery ciphertexts independently of enrolment status. See the full version [48] for further details.

IMPACT. With possession of the user key, the adversary is able to read and modify all vault data of the target user.

REQUIREMENTS. The user logs in to LastPass in the browser extension client.⁷

4.2 Item-Level Vault Encryption

As we have mentioned, even though the user interacts with the abstraction of a “secure vault”, the individual items and the user settings that comprise a vault are managed and (typically) encrypted independently. Hence, vaults are not monolithic cryptographic objects.

If the vault were monolithic, vault integrity would be easy to achieve by using authenticated encryption (AE). However, every minor change would necessitate re-encrypting the entire vault, making its maintenance and synchronisation with the server costly. Thus, a common solution is instead to encrypt each item separately with AE under a single key – but this does not necessarily result in integrity guarantees for the whole vault. For example, items could be swapped around or removed entirely without the client being able to detect it.

The attacks that follow all exploit the lack of *vault integrity* for the vendors we study. Not only can the adversary swap and delete items, they can also change metadata, settings such

⁵ This feature is only available in the paid versions of LastPass [40], but, as with Bitwarden, the code implementing this feature is present in the client independent of their subscription status.

⁶ This process is described as a “Key Exchange” by LastPass [41]. This does not correspond to the conventional meaning of the term.

⁷ Only the browser extension client supports the “key exchange” procedure used for password reset, thus the attack works only against this client.

as the number of KDF iterations, or the organisations to which a user belongs. For space reasons, we defer the description of Attacks [BW04](#), [LP05](#) and [LP06](#) to the full version [48].

4.2.1 Bitwarden

In Bitwarden, items and user settings are not uniformly protected. Rather, vault items are encrypted individually, and many settings are not encrypted at all. This also holds true for items in organisations and organisation-wide settings.

Each item is an object with various fields, such as item name, password, username, and URLs. All of these fields are individually encrypted using AES-CBC-HMAC. Up until version 2024.2.0 of Bitwarden, the fields of all items are encrypted directly with the user key k_u . As of v. 2024.2.0, “per-item keys” are used instead: a fresh key k_u^i is sampled and used for each new vault item. The item key is itself encrypted with the user key k_u , and added as a field to the item. Note that old items are not migrated to this new format.

The client retrieves vault items from the server via a sync request; the server will return the item encoded in a JSON object, with the (unencrypted) field names as keys, and the respective ciphertexts as values.

BW05: Item Field Swapping. Nothing binds a ciphertext to a particular item field. Hence, an adversary can swap around ciphertexts from different fields within an item.

For all items created before the per-item keys feature was introduced, an adversary can even swap ciphertexts around among different items, as well as create new items by merging fields from other items.

This attack has no requirements, and represents a clear violation of vault integrity. It also functions as a building block for other attacks.

BW06: Icon URL Item Decryption. Items can include a URL field, which is used to autofill the credentials and display an icon on the client. The client decrypts the URL and fetches the icon from the server, including in its request the domain and top-level domain of the URL. For instance, if the URL is “https://host.tld/path”, the client request includes “host.tld”.

This means that the adversary can learn (part of) the contents of URL fields. Using Attack [BW05](#), an adversary can place the ciphertext of sensitive item fields, such as a username or a password, in the encrypted URL field. After fetching the item, the client will then decrypt the ciphertext, confusing it for a URL. If the plaintext satisfies some conditions (i.e. containing a ‘.’ and no !), it will be leaked to the adversary.

A URL checksum feature was deployed in July 2024, making the clients store a hash of the URL in another encrypted item field, therefore providing a rudimentary integrity check and preventing this attack. Note that old items are never updated to add such a checksum: this feature only protects items created after its introduction. Furthermore, URL checksums are only checked if a per-item key is present for the item.

As we will see, an adversary can prevent per-item keys from being enabled with Attack [BW10](#).

IMPACT. The adversary can recover selected target ciphertexts in the item, such as the username or the password.

REQUIREMENTS. The user opens a vault containing items that do not use per-item keys (i.e., items created before July 2024, or after Attack [BW10](#) is run). The target plaintext must satisfy some additional conditions, detailed in the full version [48].

BW07: Remove KDF Iterations. Client authentication to Bitwarden is performed by checking if a hash of the master key h_m , derived from the master password input by the user at login time, matches a hash stored on Bitwarden’s servers. In order to protect the master password from brute-force attacks, computation-intensive Password-Based Key Derivation Functions (PBKDFs) are used to compute this hash. The derivation speed is determined by the parameters of the PBKDF.

For instance, when PBKDF2 is used with the default 600,000 iterations, the hash is obtained as $k_m \leftarrow \text{PBKDF2}(pw, email, 600000)$; $h_m \leftarrow \text{PBKDF2}(k_m, pw, 1)$.

The KDF settings are stored unauthenticated and unencrypted on the server. When a target user logs in, the adversary can choose an arbitrary PBKDF configuration, including one with a minimal iteration count. The user will then send a master key hash that is efficient to compute during authentication. In our example, this means that the adversary can set the first iteration count to 1, thereby reducing the number of PBKDF2 iterations from 600,001 to 2 and achieving a 300,000x speed-up in a brute-force attack.

IMPACT. The adversary can reduce the cost of a brute-force attack on a target user’s master password; successfully recovering the password enables full control of the user’s account.

REQUIREMENTS. The user logs in.

4.2.2 LastPass

Similarly to Bitwarden, vault items in LastPass are composed of various fields such as item name, username, URL, and password, and all of these fields are individually encrypted under the same user key k_u . Similar attacks therefore apply.

LP02: Item Field Swapping. There is no associated data binding each ciphertext to its field. Hence, an adversary can swap around ciphertexts from different fields within an item, as well as between different items, and create new items by merging fields from other items.

LP03: Icon URL Item Decryption. LastPass login items can include a URL. This is used to autofill the credentials and display an icon on the client. In order to fetch these icons, the client makes a query to LastPass’s icon server, including in the request the URL that was stored inside the item. Just like in Bitwarden, an adversary can use the item-field swapping attack (Attack [LP02](#)) and place the ciphertext of sensitive

item fields, such as the username or the password, in the URL field. The plaintext will be leaked to the adversary via the query to the icon server. Note that, unlike Bitwarden, there are **no conditions** on the plaintext formatting.

IMPACT. The adversary can recover (parts of) selected target ciphertexts in the item, such as the username or the password.

REQUIREMENTS. The client opens the modified vault.

LP04: Remove KDF Iterations. Client authentication to LastPass is performed by checking if a hash of the user key – derived from the master password input by the user at login time – matches a hash stored on LastPass’s servers. Just like in Bitwarden, 2 invocation of PBKDFs are used to derive this hash, with a default of $600,000 + 1$ iterations for PBKDF2.

Similarly to Bitwarden, the KDF parameters are unencrypted and unauthenticated, enabling an adversary to choose them arbitrarily. The PBKDF2 iteration count can then be lowered down to 3 (a client-enforced minimum of 2 for the inner invocation, while the outer is fixed to 1). The adversary thus achieves a speed-up of 200,000x in a brute-force attack.

IMPACT. The adversary can reduce the cost of a brute-force attack on a target user’s master password; successfully recovering the password enables full control of the user’s account.

REQUIREMENTS. The user logs in.

4.2.3 Dashlane

Recall that items in Dashlane are treated as transactions and encrypted with a key derived from the master key based on transaction parameters. For all transactions sharing the same parameters (in particular, KDF and cipher mode; the salt is constant for all transactions) the same derived key is used.

DL01: Transaction Replay. Since the transaction key does not depend uniquely on the transaction, the server can duplicate, reorder or drop transactions without the client noticing.

This attack trivially violates vault integrity. It also serves as a building block for other attacks, see Section 4.4.2.

4.3 Sharing

These attacks target the *sharing* of items. All of the password managers we study implement some kind of sharing, and they fall victim to attacks due to a lack of user authentication.

The attacks in this category range in severity: while the vulnerabilities in LastPass and Dashlane only affect the items being shared, the ones in Bitwarden enable some of our Key Escrow attacks and have further-reaching consequences.

4.3.1 Bitwarden

Recall that in Bitwarden, users can share items with other users in organisations; each organisation has a symmetric key k_{org} , which is stored on the server encrypted under the public RSA key of each member of the organisation.

BW08: Organisation Injection. There is no cryptographic binding between a user’s vault (or even the vault abstraction) and the organisations a user is part of. Users learn the list of the organisations they are part of when their client performs a sync to fetch the latest data from the server. Furthermore, copies of the encrypted organisation symmetric key are not authenticated in any way: anyone can encrypt a chosen symmetric key under the public key of any user.

It follows that adding a user to an organisation only requires access to their public key, which is known by the server. Thus, an adversary can create a new organisation and trivially add users to it: they just need to encrypt the organisation symmetric key k_{org} under the target user’s public key. The client will silently accept the new organisation at the next sync.

As a corollary, the adversary can also add arbitrary users to any organisation whose symmetric key they know. We give the full details of the attack in the full version [48].

IMPACT. The adversary can add users to arbitrary organisations, thus violating vault integrity by adding organisation items to the user’s vault. This can be used in framing attacks, for example, by planting incriminating material. Furthermore, it enables several other attacks (e.g., BW03 and BW09).

REQUIREMENTS. The client performs a sync operation. This is done automatically by the client at startup, when the vault page is loaded, in response to some user actions, and upon server-generated events.

BW09: Organisation Overwrite. Immediately after an organisation is created, the client of the creator performs a sync to fetch the latest data from the server. As part of this, the keys of the newly created organisation are fetched. After the sync is completed, the creator of the organisation is able to decrypt the organisation symmetric key using their private RSA key and use it to decrypt all organisation data.

By applying the same method as in Attack BW08 above, the adversary can overwrite the newly created organisation: they can swap the encrypted symmetric key c_{org} and the organisation data with values they control. Upon sync, the client has no way of distinguishing such a malicious key from the legitimate organisation key.

In contrast to Attack BW08 (which creates a new organisation), this attack is entirely undetectable. Note that it is also possible for an adversary to overwrite an organisation after it is created, but in doing so they would delete any pre-existing organisation items.

IMPACT. The adversary can read and modify all the items in the newly created organisation.

REQUIREMENTS. The user creates an organisation.

4.3.2 LastPass

Recall that LastPass implements item sharing by means of shared folders. Items in a shared folder are encrypted with

the symmetric shared folder key, which is in turn encrypted under the public key(s) of the recipient(s) of the folder.

LP07: Sharing Key Overwrite. During an item-sharing operation initiated by a user, the server, acting as an adversary, sends an adversary-controlled recipient RSA public key instead of the intended recipient key.

The client (and the user) cannot distinguish the malicious key from the intended one: the recipient RSA public keys are not authenticated in any way, the user interface does not allow for out-of-band verification of key fingerprints, and no key transparency is implemented.

IMPACT. The adversary can read and modify all the items in a shared folder.

REQUIREMENTS. The user shares an item with any recipient.

4.3.3 Dashlane

Like LastPass, Dashlane implements item sharing using public key encryption. The problem of unauthenticated public keys described in Attack LP07 also affects Dashlane.

When a client shares an item with a user or in a group, a fresh symmetric sharing key (k_{shr}) is sampled. This key is either used to directly encrypt the shared item (in the case of user-to-user sharing), or to encrypt group keys which in turn protect shared items (for group sharing). Regardless of setting, the sharing client retrieves the RSA public keys of the recipient(s) from the server and, without verifying their authenticity, encrypts k_{shr} under each recipient's public key, yielding c_{shr}^{user} for each recipient $user$. The client sends each c_{shr}^{user} to the server for storage.

DL02: Sharing Key Overwrite. When the sharing client requests a recipient public key, an adversary can overwrite it with an RSA key pk^{adv} it controls, causing k_{shr} to be encrypted under the malicious public key. The resulting ciphertext c_{shr}^{adv} is then sent to the adversary, allowing them to decrypt it with sk^{adv} and recover k_{shr} . This in turn gives the adversary access to the shared item(s) protected by k_{shr} .

IMPACT. The adversary can read and modify all shared items.

REQUIREMENTS. The user shares an item with any recipient.

4.4 Backwards Compatibility

These attacks exploit backwards compatibility features of Bitwarden and Dashlane. Both password managers have evolved the cryptographic primitives they support over the years, shifting from AES in CBC mode for encryption and PBKDF2 for key derivation to AE schemes such as AES-CBC-HMAC and memory hard PBKDFs like Argon2d. Both also show changes in the key derivation trees, with Bitwarden migrating to better key separation with the introduction of per-item keys.

The current client versions only use modern cryptography by default, but they still offer “legacy” support. The attacks described below enable the adversary to downgrade, sometimes

permanently, the primitives used by the clients, reducing the security of the affected password managers to the security of the weakest supported cryptographic algorithm.

4.4.1 Bitwarden

The latest Bitwarden clients use authenticated encryption (AE) for all symmetric encryption operations by default, thereby ensuring both confidentiality and integrity of the encrypted data. The AE scheme used is AES-CBC with HMAC, combined via the encrypt-then-MAC paradigm. In particular, this encryption scheme is used to decrypt the user key received from the server at login time and to encrypt and decrypt the items inside the vault.

However, older versions of Bitwarden used AES-CBC (without HMAC), which does not offer any integrity protection. To accommodate both formats, each ciphertext has an encryption type attribute (type 0 for AE, type 1 for plain AES-CBC). The encryption type is prepended to the ciphertext and is not integrity-protected.

Furthermore, older versions of Bitwarden did not use the current key hierarchy. To support the old hierarchy, a modern client will generate a new RSA keypair for the user (pk', sk') if one was not provided by the server, and use k_m to encrypt sk' if no user key ciphertext c_u was provided by the server. Both of these features can be seen as backwards compatibility measures, allowing a modern client to support old user vaults and to migrate the user vault to the current feature set.

BWb1: Known Plaintext Oracle Gadget. An adversary can leverage the generation of RSA keypairs and the support for encryption using the master key k_m to obtain AES-CBC encryptions of known plaintexts under an (unknown) k_m .

At login time, the adversary responds to the login request with a null c_u and a null private RSA key ciphertext c_{sk} . The absence of c_{sk} will force the user to generate a new RSA keypair. However, since c_u was also set to null, the client will use k_m to encrypt the new secret key using AES-CBC, yielding c'_{sk} , and will then send pk', c'_{sk} to the server.

Notably, sk' also contains an encoding of pk' . Thus, by observing the encrypted c'_{sk} and pk' , the adversary learns a 5-block-long plaintext-ciphertext pair.

BW11: User Key Overwriting. The adversary can forge a ciphertext c_u^{adv} , such that the client will decrypt it to a user key $k_u = k_u^{enc} || k_u^{mac}$ for which the adversary knows k_u^{enc} . Recall that k_u^{enc} and k_u^{mac} are 32 bytes each, spanning two blocks of CBC plaintext.

The adversary constructs c_u^{adv} by using the 5-block plaintext-ciphertext pair from BWb1. Via bit-flipping, the adversary sets the last block of ciphertext to a block P_5 of padding, obtaining a garbled block P_4 , and known blocks P_1, P_2, P_3 . Finally, the adversary marks this ciphertext as type 0.

Upon receiving c_u^{adv} , the client will accept it as a valid user key encryption: the ciphertext is encrypted under the

master key k_m , the type is 0 so CBC decryption is used, and unpadding succeeds. The client then parses $k_u = k_u^{enc} \parallel k_u^{mac}$. Since k_u^{enc} consists of blocks P_1 and P_2 , the adversary knows k_u^{enc} fully, while k_u^{mac} is unknown to the adversary because of the garbled block P_4 .

IMPACT. The adversary learns k_u^{enc} , used to encrypt all the vault items. Therefore, the confidentiality of all items encrypted after the attack is mounted is forfeit. Integrity is preserved, as the adversary does not learn k_u^{mac} .

REQUIREMENTS. The user logs in. Due to the attack, this initial login attempt fails, and the user is redirected to a migration page. Regardless of the action taken by the user after the migration page is displayed, the attack completes successfully during a subsequent login attempt.

BW12: Downgrade to Legacy. The adversary can forge a ciphertext c_u^{adv} , such that the client decrypts it to a user key $k_u = k_u^{enc} \parallel k_u^{mac}$, where k_u^{mac} is empty.

This can be achieved by following the same steps as in Attack [BW11](#), but using only three blocks of ciphertext: in this case, P_3 is set to a block of padding, and P_2 is garbled. As previously discussed, the client will accept this c_u^{adv} as valid. The ciphertext being too short triggers another backwards compatibility feature: the client parses $k_u = k_u^{enc}$ without the “mac” key part, causing all subsequent encryptions to be performed using CBC-only mode. Note that if per-item keys are enabled, item encryption still uses AES-CBC with HMAC. But, as we show in [BW10](#) (see full version [48]), the adversary can easily downgrade the client to not use per-item keys.

IMPACT. The attack downgrades encryption to CBC-only mode, trivially compromising the integrity of all items encrypted after the attack is executed. In turn, this enables a padding oracle attack [54], leading to a complete compromise of vault confidentiality. We did not develop a PoC for this last aspect of the attack, but the developers of Bitwarden acknowledged its feasibility. We also note that the adversary knows P_1 , and therefore learns half of the encryption key k_u^{enc} .

REQUIREMENTS. The same as [BW11](#).

4.4.2 Dashlane

Recall from Section 3.3 that every transaction has a *payload*, which defines the key derivation function, cipher, and mode to use to decrypt the transaction content. Two payload formats are relevant for our attacks: KWC3, using PBKDF2-SHA1 and AES-256-CBC; and Flexible, with a format string which is parsed to select the KDF (PBKDF2-SHA256 or SHA1, Argon2i, or the default Argon2d), its parameters, the cipher (only AES-256), and the cipher mode (the default CBC-HMAC and GCM being the only documented modes).

Modern Dashlane clients default to Flexible mode; upon detecting a KWC3 payload they automatically migrate the transaction to Flexible. Support for the legacy format varies across clients: we verified that the web extension supports it,

while the CLI client does not. All of the following attacks only affect clients that support the legacy format. We produced PoCs against the web extension client.

DLb1&2 and DL03: Padding Oracle and Injection. Unexpectedly, CBC-only encryption is still accepted as an encryption mode in the Flexible payload format and does not trigger a migration. This enables a padding oracle attack [54], which the adversary can use to decrypt ciphertext under key k_{cbc} (DLb1), or encrypt arbitrary plaintext under k_{cbc} [29, §3] (DLb2). Our unoptimized attack runs at a rate of 14 days per decryption or encryption block. For space reasons, we defer the full description of these attack gadgets to the full version [48]. Attack DL03 (Item Injection), which uses DLb2 to forge arbitrary transactions, is similarly deferred to [48].

DL04: Remove KDF Iterations. User settings in Dashlane control many client behaviours. Among these, they define what KDF to use for all new transaction payloads, and the respective parameters.

Using Attack DL01, the adversary can remove any existing transaction of type Setting. This moves the user settings to a known clean state. Then the adversary uses Attack DL03 to forge a new transaction of type Setting. When the user creates a new item, the client will derive a key from the master password with the adversary-chosen KDF parameters.

For instance, the adversary can change the number of PBKDF2 iterations from the default of 200,000 iterations to 1, achieving a 200,000x speed-up in a brute-force attack.

IMPACT. The adversary reduces the cost of a brute-force attack on a target user’s master password; success enables full control of the target user account.

REQUIREMENTS. Execution of DL01 and DL03.

DL05: CBC-Only Downgrade. The setting transaction also controls which cipher and mode of operation the client uses for future transactions. Similarly to Attack DL04, the adversary can forge settings instructing the client to use CBC-only encryption. When the user creates new items, the client will then always use CBC-only mode. The resulting transaction will be vulnerable to the padding oracle attack in DLb1.

IMPACT. The adversary can decrypt all the transactions produced after the attack is completed.

REQUIREMENTS. Execution of Attacks DL01 and DL03, plus one execution of DLb1 for each decrypted transaction.

Attack timing. In our PoCs, we successfully forge a minimal transaction of type Setting which is only 9 blocks long, enabling us to test Attack DL04 and DL05. We artificially sped up the execution for testing, but we estimate this would take more than 126 days at the normal attack speed. This is a long time, but it is reasonable for a targeted attack.

For space reasons, we defer the description of one further attack (DL06) in a similar style to the full version [48].

5 Mitigations

Our attacks can all be mitigated using a combination of authenticated encryption (AE), key separation (KS), plaintext authentication (AD), public key authentication (PKA), and ciphertext authentication (SC). We detail how in Section 5.1; Table 1 additionally maps attacks to mitigations.

In general, these techniques can be used to design the core of a secure E2EE password manager: in Section 5.2 we briefly sketch such a high-level design.

5.1 Immediate Remediations

In our ongoing discussions with the vendors, the introduction of breaking changes emerges as the major obstacle to deploying our countermeasures. As we argue in Section 6, breaking changes are necessary, and even inevitable. To this aim, we propose the use of specialized password manager clients, with no functionality besides implementing a forced migration to the new vault format. This would prevent any user from losing access to their data, while preserving security for the entire user base. Below, we detail the changes we suggest.

Only use AE (AE). Disabling legacy support for encryption schemes that provide no integrity immediately remediates most of our Backwards Compatibility attacks on Bitwarden (BWb1, BW11, BW12) and Dashlane (DLb1, DLb2, DL03, DL05, DL06), since they exploit the support of AES-CBC. Encrypting item metadata with AE would remediate BW04 (Unprotected Item Metadata). LastPass would need to introduce AE from scratch. This would remediate many of its Vault Encryption issues (LP05, LP06), and, together with the KS mitigation (see below), also attacks LP02 (Item Field Swapping) and LP03 (Icon URL Item Decryption).

Additionally using authenticated data (AD) to protect metadata fields in the items and user settings in the vault would mitigate the Vault Integrity issues in Bitwarden (BW05, BW04, BW06), LastPass (LP02, LP03, LP06), and Dashlane (DL01).

Full key separation (KS). Most vault integrity issues can be mitigated by using proper key separation: each key should be used only to encrypt a single field of a single vault item, with all keys being created from the master key using a suitable key derivation function and per-item and per-field context strings. Recall that, in contrast to encrypting the whole vault as a monolith, this approach allows efficient synchronisation of changes: each item can be encrypted and synchronised independently, while still protecting the overall integrity.

In Bitwarden, item keys provide a good starting point. Extending this approach to the field level would both make URL checksums redundant, and remediate BW05 (Item Field Swapping) and BW06 (Icon URL Item Decryption). Item keys should also be made mandatory, mitigating BW10 (Disable Per-Item Keys). Together with the AE mitigation, KS would

also prevent issues like BWb1 (Known Plaintext Oracle Gadget) from arising. LastPass currently lacks any form of key separation within the vault. Introducing it along with an AE scheme would mitigate LP02 and LP03. In Dashlane, each transaction could be encrypted with a different key, derived according to a unique transaction identifier, to resolve DL01.

Authenticate all plaintext (AD). All settings and metadata that cannot be encrypted should at least be authenticated (i.e., integrity protected). This can be achieved using Authenticated Encryption with Authenticated Data (AEAD) as an extension of AE. This would mitigate the metadata malleability issues (BW04, LP06). Further, authenticating security-critical user settings like PBKDF parameters (such as the iteration count) would mitigate the KDF attacks (BW07, LP04). The client can use the server-provided KDF parameters to derive the authentication key, use it to verify the integrity of the parameters themselves, and – in case of a mismatch – abort before any further communication with the server. On the one hand, this thwarts our immediate attacks: an adversary tampering with the KDF parameters would need to predict the authentication key for the client to accept the modified parameters. On the other hand, as pointed out in [47, §3.2], the security in this setting is circular and relies on non-standard properties of MACs. This measure should therefore be combined with allowlisting of safe sets of KDF parameters in clients.

Similarly, authenticating the last known server version seen by the client by adding it to the vault additionally mitigates BW10 (Disable Per-Item Keys). An attacker can still replay an old version of the vault (with the corresponding server version), but all newly created items are protected.

Public key authentication (PKA). Introducing proper authentication of public keys is non-trivial: a stop-gap mitigation would be storing known public keys in the vault after an (optional) out-of-band verification. This would allow the users to detect public key changes, preventing active attacks. Even when no out-of-band verification is performed, this would still prevent some attack scenarios: the attacker can then only substitute public keys on their first use.

An alternative is to introduce a PKI where a Certificate Authority (CA) signs all public keys. This CA should not be under the control of the same entity that controls the password manager's server, so that it can truly act as a trusted third party. This is particularly relevant in a business setting, given that many companies that use password managers may already operate their own CA infrastructure.

These measures would remediate most of the Sharing issues (BW09, LP07, DL02). They would also allow the user to authenticate the public keys of admins and organisations used for Key Escrow. Additionally, assuming that a correctly signed list of admins is published, this would remediate several Key Escrow issues (BW01, BW03, LP01).

Ciphertext authentication (SC). In Bitwarden, the only cryptographic check of organisation membership is through the decryption of organisation data, encrypted by the user under their own public key. However, as we have noted, anyone can produce such ciphertexts. This results in [BW02](#) (Malicious Key Rotation), [BW08](#) (Organisation Injection), and [BW03](#) (Malicious KC Conversion). These attacks can be prevented by switching to signcryption [56]. By signing plaintexts before encryption, clients can verify that the user themselves is the sender. An alternative would be to forgo public-key encryption entirely in favour of symmetric (authenticated) encryption. This is possible since the user is both the sender and the recipient of the affected data. An added benefit of this approach is efficiency gains from removing unnecessary public-key cryptography. However, it might require a larger re-design compared to simply adding signatures. Either approach mitigates all of the above attacks.

A combination of signcryption and [PKA](#) is needed when receiving public key ciphertext encrypted by other users, such as organisation invites in Bitwarden, or the list of admins for a team in LastPass. This would be needed to remediate [BW01](#) (Malicious Auto-Enrolment), [BW09](#) (Organisation Overwrite) and [LP01](#) (Malicious Password Reset).

5.2 Future Designs and Limitations

The mitigations we propose can serve as a roadmap for a provably secure password manager design, with minimal performance overhead compared to existing designs, as follows. Starting from a PBKDF-derived key computed from the user’s master password, standardized and efficient key derivation functions are used to derive separate symmetric keys for all vault items, thus providing proper key separation. With these keys, vault data is encrypted using a modern AEAD scheme. Using symmetric cryptography with built-in authentication both provides ciphertext integrity and has low performance cost compared to using bespoke mechanisms for achieving vault integrity. All user data is both encrypted and authenticated, with only unavoidable metadata, such as PBKDF parameters, left as authenticated-only additional data. Wherever public-key cryptography is necessary (e.g., for sharing), measures such as certificates are taken to authenticate public keys, and signatures are used to authenticate ciphertexts.

Caveat. We have presented mitigations for all the attacks we found, but there may be other cryptographic attacks yet to be discovered for which our mitigations would be ineffective. Formally defining the security goal of a password manager in the face of a malicious server is a necessary first step toward provable security. The same applies to our high-level sketch of a secure design: only a security analysis backed by formal models and proofs can properly inform a more detailed construction; such analysis is beyond the scope of this paper.

6 Discussion

Password managers deal with critically sensitive data – yet, we have exposed severe attacks against three important vendors in this space. This section explores the manifold reasons for those attacks, discusses broader impacts of our findings, and provides directions for future work.

A burdensome legacy. Password managers were early entrants in the world of end-to-end encrypted applications: they predate all modern end-to-end encrypted messaging applications, with the initial release of LastPass dating back to 2008. Encryption of vaults has always been a focus of these products, but cryptographic best practices have changed a lot in the last years. Dashlane and Bitwarden are more recent, with first releases in 2012 and 2016 respectively, but they were entering a market with quite low cryptographic standards. The early releases of both show a taste for early-2000s cryptography, with lack of integrity protection for encryption and unprincipled designs being prevalent.

Motivated by backwards compatibility, residues of this legacy cryptography have remained across codebase changes and redesigns. In our communication with vendors, they expressed deep concern about leaving users with impossible-to-decrypt vaults because of cryptographic changes. This motivated the extreme lengths to which some of them go in order to support old formats. In the meantime, the security community had learned, through the long saga of attacks on SSL and TLS [2, 6, 7, 13–16, 32, 53], to put aside backwards compatibility and orient itself towards *cryptographic alacrity*: introducing sharp, but clean, breaking changes to keep the ghosts of past cryptographic attacks at bay.

Deceptively simple. Unlike secure channel protocols and end-to-end encrypted messaging, password managers have arguably escaped deep academic scrutiny before now.⁸ Perhaps this is because the problem they address looks, on the surface, like a simple matter of key derivation and then encryption, without any interesting problems to solve. After a closer look, we have seen that password managers are far from simple: they have evolved to include complicated protocols for key synchronisation, recovery and rotation, sharing of encrypted elements, and migration between different cryptographic primitives. The complexity they have reached is comparable to that of modern E2EE cloud storage systems.

Even though password managers have evolved, with cryptographic best practices trickling in over time (e.g. AEAD and memory-hard PBKDFs being eventually integrated in Dashlane and Bitwarden), there was never a gold standard to match (like TLS or Signal), or even precise security notions to aim for. This is similar to the situation for E2EE cloud storage, where a recent rash of research [3, 4, 11, 35, 37] has exposed the frailties of deployed systems, and formal

⁸ They are frequently subjected to code audits and security reviews, but not to *cryptographic* reviews.

foundations closely linked to practice are only now emerging. Indeed, the two applications are close enough that the recent foundational work of [10] could provide a good starting point for password managers too.

Integrity in motion. To draw another parallel with cloud storage, password managers also face the challenge of efficiently synchronizing encrypted data. The naïve approach of storing all user files in a single AE-protected blob is clearly unviable for large amounts of data. Modern password managers, supporting large numbers of items, consequently structure their vaults to enable individual item synchronisation.

As we suggest in Section 5, the apparent tension between integrity and incremental updates can be solved by clean key separation. Preventing deletion remains an open problem; Merkle tree techniques can prevent selective deletion attacks by authenticating the vault structure, but complete roll-back of vaults cannot be ruled out without outsourcing versioning information to a trusted third party or local cryptographic state. The latter approach is problematic for password managers, since users expect access to their vault upon presenting only their master password, even on a stateless client.

Share with care. Many of the vulnerabilities we identify rely, in one form or another, on substituting public keys with adversary-controlled ones. Such attacks have been known since the dawn of public key cryptography [27, §III].

When it comes to verifying public keys of other users, administrators or organisations, password managers run into a very well-known problem: there is no way around the need to provide some form of verification of public keys. In short, PKI (of some form) cannot be avoided.

In our mitigations, we propose relying on PKIs or on Trust-On-First-Use (TOFU) systems. Secure messaging applications are currently exploring the use of Verified Key Directories (VKDs) to overcome this problem. While VKDs would not prevent a malicious server from associating the wrong public key with a user, they make this kind of tampering evident, so that users are alerted when their providers misbehave.

Authenticating public keys is only half the battle. It seems to be a common misunderstanding both here and for E2EE cloud storage [3, 37] that public key encryption (PKE) somehow produces authenticated ciphertexts. In fact, *no* PKE scheme can offer data origin authentication (because, by definition, any party can encrypt). This is why we suggest the usage of primitives like signcryption [56] to correctly bind ciphertexts to their senders.

Securing a “backdoor”. Password managers attempt to implement mechanisms for recovery of the vault content in the case the user forgets their master password. There is a clear tension here between security and usability: not having any recovery mechanism provides for the best possible security, but might be untenable for large populations of non-technical users. Admin-assisted recovery can be implemented

safely, but success hinges on the key authentication problem discussed above being solved first. Again, we can look at the secure messaging space for state-of-the-art solutions to this problem: messengers like Signal and WhatsApp [22, 25] are converging to the usage of hardware security modules (HSM) in the cloud for protecting account recovery and backups. Similarly, Apple’s Advanced Data Protection uses HSMs to allow user to recover their key material from just their PIN [8].

Implementing part of the Key Escrow functionality on an HSM can offer at least some form of guarantee that the server is behaving honestly. Hybrid solutions could include using the HSMs to distribute the keys of the admin, without having to rely on other forms of PKI or out-of-band verification. This is an interesting future direction in the context of formal security guarantees for password managers.

7 Conclusions

We have argued that the malicious server threat model is the correct one for cloud-based password managers: security in this setting is both claimed in vendor advertisement and represents the gold standard for modern end-to-end encrypted systems. We then showed that three popular password managers with more than 60 million users in total are badly broken in this threat model. The vulnerabilities that we describe are numerous but mostly not deep in a technical sense. Yet they were apparently not found before, despite more than a decade of academic research on password managers and the existence of multiple audits of the three products we studied. This motivates further work, both in theory and in practice.

Beyond showcasing the relevance of the malicious server threat model in practice, our attacks show that these systems would benefit from theoretically well-founded designs. Vendors need to update their products to use modern cryptographic primitives and best practices – but to ensure solid foundations, novel definitions to capture security in this setting are also needed. Furthermore, the common issues with features such as key escrow, item-level vault encryption, password sharing, and backward compatibility represent interesting questions for the academic community on how to achieve security while still providing these functionalities.

Working together with vendors to fix the vulnerabilities, our research has contributed to improved security in practice for millions of affected users. It has also increased the vendor-driven interest in formal guarantees of security. We hope that this paper serves to help raise the relevant industry sector’s knowledge about encryption further above zero. At the same time, we hope our work spurs the academic community to further research in this area – as well as to consider what more it can do to convert its research artefacts into meaningful impacts on industrial practice.

A Ethical Considerations

Our work concerns the security analysis of three products claimed to have tens of millions of users in a threat model that we have argued as being worthy of serious consideration. Our analysis has shown these products to be badly broken in our chosen model. This situation certainly calls for a detailed statement concerning possible harms arising from our research, how these are mitigated in our approach, and the extent to which they are balanced by other outcomes. We formulate this in terms of a stakeholder-based analysis. We take a consequentialist viewpoint to conduct our analysis.

The main stakeholders related to our work, as we see them, are: us, the researchers; the vendors of the affected products (Bitwarden, Dashlane, LastPass); vendors of related products; customers using the affected products (both individuals and organisations); society at large.

We eschew further discussion of ourselves as stakeholders, since we cannot readily set aside our own feelings on the matter to provide a rational analysis of how we might be affected by the publication (or otherwise) of our work.

Concerning the vendors of the affected products: they may suffer reputation and thence commercial damage as a result of our work being published. For example, it has been reported that LastPass has lost many customers due to repeated breaches [34, 49]. This potential damage is mitigated in our approach by providing the affected vendors with a 90-day disclosure window during which we did not talk in public about the research. During this window, we supported the vendors, initially with very detailed descriptions of the security issues we have found, then repeatedly offering support to address the issues (video-conferences, review of patches). We believed that 90 days should be sufficient time to address the identified issues, but we always remained open to extending this time period upon request, and actually did so for one vendor. We consider this approach to be relatively standard for the field of software vulnerability disclosure, except perhaps for the level of our engagement in the triage process. For example it is endorsed by ENISA, an EU body.⁹

We disclosed to Bitwarden on 27.01.2025, to LastPass on 04.06.2025, and to Dashlane on 29.08.2025. We had a video-conference and numerous email exchanges with Bitwarden. At the time of writing, they are well advanced in deploying mitigations for our attacks: **BW01**, **BW03**, **BW11**, **BW12** were addressed, the minimum KDF iteration count for **BW07** is now 5000, and their roadmap includes completely removing CBC-only encryption, enforcing per-item keys and changing the vault format for integrity. On 22.12.25 they shared with us a draft for a signed organisation membership scheme, which would resolve **BW08** and **BW09**. At our request, to maintain anonymity, they have not yet credited us publicly for the disclosure, but plan to do so. In the case of LastPass, our

initial contact via the official e-mail channel on 04.06.2025 did not receive a response. One of us reached out to the CTO over social media, at which point our original e-mail was found in the organisation's spam filter on 26.06.2025. We mutually agreed with LastPass to reset the 90-day clock. They provided us with regular updates and we also met for a video-conference to clarify our findings and discuss their mitigation plans. They accepted our request to work with them directly rather than with their preferred bug bounty program because the latter would impose publication restrictions that were not acceptable to us; on 7.10.25 they awarded us two bug bounty payments. At the time of writing, **LP03** has been addressed. Our initial email to Dashlane was also lost, but our second attempt at contacting them on 05.09.2025 was successful. Since then, we worked with their engineers to replicate our attacks, and had several discussions over email as well as one video-conferencing call. At the time of writing, **DL03**, **DL04**, **DL05** and **DL06** have been mitigated by disallowing CBC-only mode in Flexible payloads; no remediation is planned for **DL01** and **DL02**.

Upon agreement with all involved vendors, we shared a preprint of our research paper with them in a joint thread, allowing them to also discuss the findings and take-aways among themselves. We also coordinated with them on holding back on public statements concerning the vulnerabilities until the mutually agreed date of 14.02.2026, significantly later than the end of any of the 90-day disclosure periods.

Concerning vendors of related products: there is a clear risk that the issues we have found in the three vendors we have studied may extend to other vendors, either *mutatis mutandis* or with more substantial yet simple changes. In mitigation to the possible harms, we are aware that there is an informal industry forum where our work was already discussed as a result of the disclosure to Bitwarden. While this formally broke any agreement concerning coordinated disclosure (and arguably relieved us of the burden to maintain confidentiality), we also found it understandable and saw it as having a positive effect in alerting other vendors to the possibility of problems in their products. Vendors who have a mature security mindset could already start to analyse their own systems. We will also publicise our work widely after the end of the disclosure periods, including by sharing it (by request from Dashlane) with a group of CTOs of other password managers on the market.

We turn now to customers of this collection of vendors. Although we consider the malicious threat model to be realistic, we also consider that the majority of users would not be targeted by attacks of the type we have presented because of the stringent conditions needed to mount them. Higher risk individuals and organisations, however, could be at substantial risk of harm. Unfortunately, we cannot exclude the possibility that our attacks were already known to advanced threat actors – after all, we have learned from the Snowden revelations that national security agencies are routinely tasked with penetrating systems like the ones we analyse and are willing

⁹ <https://www.enisa.europa.eu/publications/coordinated-vulnerability-disclosure-policies-in-the-eu>

to conduct active attacks on targets. The best mitigation for these parties is to trust that vendors will rapidly and effectively patch their systems, and here we have made real effort to engage with the affected vendors to assist them in this process.

Finally, we see society at large as benefitting from our work, in that our work should help to raise the security bar for password managers, pushing their vendors to either improve security or make clearer statements about what security their systems actually provide, so that customers can judge (perhaps with the help of expert guides) whether the products meet their requirements or not. In the longer term, we see our work as contributing to reducing the possible harms from using password managers by making them more secure.

Given the overall balance of benefits versus harms, and the extensive mitigations to limit those harms, we took the decision to proceed with the research and to publish it.

B Open Science

Our artefacts consist of proof-of-concept code for all the attacks in the paper. While some of the attacks are in the process of being mitigated, others remain exploitable due to the need for longer-term mitigations (e.g. public key authentication). Due to the sensitivity of the PoCs, we came to the decision that we should also not make these PoCs generally available before all of the agreed embargo periods were complete. However, we made all the artefacts available to the program committee upon request. At the precise time of writing, all the 90-day disclosure periods have ended, but we have agreed a common public disclosure date with the affected vendors of 14.02.2026. We will make PoCs for mitigated attacks publicly available at that date at the following location: <https://doi.org/10.5281/zenodo.17977565>.

References

- [1] Harold Abelson, Ross J. Anderson, Steven M. Bellovin, Josh Benaloh, Matt Blaze, Whitfield Diffie, John Gilmore, Matthew Green, Susan Landau, Peter G. Neumann, Ronald L. Rivest, Jeffrey I. Schiller, Bruce Schneier, Michael A. Specter, and Daniel J. Weitzner. Keys under doormats. *Commun. ACM*, 58(10):24–26, 2015.
- [2] David Adrian, Karthikeyan Bhargavan, Zakir Durumeric, Pierrick Gaudry, Matthew Green, J. Alex Halderman, Nadia Heninger, Drew Springall, Emmanuel Thomé, Luke Valenta, Benjamin VanderSloot, Eric Wustrow, Santiago Zanella-Béguelin, and Paul Zimmermann. Imperfect forward secrecy: How Diffie-Hellman fails in practice. In Indrajit Ray, Ninghui Li, and Christopher Kruegel, editors, *ACM CCS 2015*, pages 5–17. ACM Press, October 2015.
- [3] Martin R. Albrecht, Matilda Backendal, Daniele Coppola, and Kenneth G. Paterson. Share with care: Breaking E2EE in nextcloud. In *2024 IEEE European Symposium on Security and Privacy*, pages 828–840. IEEE Computer Society Press, July 2024.
- [4] Martin R. Albrecht, Miro Haller, Lenka Mareková, and Kenneth G. Paterson. Caveat implementor! Key recovery attacks on MEGA. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part V*, volume 14008 of *LNCS*, pages 190–218. Springer, Cham, April 2023.
- [5] Martin R. Albrecht and Kenneth G. Paterson. Analyzing cryptography in the wild: A retrospective. *IEEE Secur. Priv.*, 22(6):12–18, 2024.
- [6] Nadhem J. AlFardan, Daniel J. Bernstein, Kenneth G. Paterson, Bertram Poettering, and Jacob C. N. Schuldt. On the security of RC4 in TLS. In Samuel T. King, editor, *USENIX Security 2013*, pages 305–320. USENIX Association, August 2013.
- [7] Nadhem J. AlFardan and Kenneth G. Paterson. Lucky thirteen: Breaking the TLS and DTLS record protocols. In *2013 IEEE Symposium on Security and Privacy*, pages 526–540. IEEE Computer Society Press, May 2013.
- [8] Apple Support. Advanced Data Protection for iCloud. <https://support.apple.com/guide/security/advanced-data-protection-for-icloud-sec973254c5f/web>, September 2024. Accessed: 2025-08-25.
- [9] Gildas Avoine, Xavier Carpent, and Diane Leblanc-Albarel. In the vault, but not safe: Exploring the threat of covert password manager providers. Cryptology ePrint Archive, Report 2025/1278, 2025.
- [10] Matilda Backendal, Hannah Davis, Felix Günther, Miro Haller, and Kenneth G. Paterson. A formal treatment of end-to-end encrypted cloud storage. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part II*, volume 14921 of *LNCS*, pages 40–74. Springer, Cham, August 2024.
- [11] Matilda Backendal, Miro Haller, and Kenneth G. Paterson. MEGA: Malleable encryption goes awry. In *2023 IEEE Symposium on Security and Privacy*, pages 146–163. IEEE Computer Society Press, May 2023.
- [12] A. Belenko and D Sklyarov. “Secure password managers” and “military-grade encryption” on smartphones: Oh, really? Technical report, Elcomsoft Co. Ltd., 2012.

- [13] Benjamin Beurdouche, Karthikeyan Bhargavan, Antoine Delignat-Lavaud, Cédric Fournet, Markulf Kohlweiss, Alfredo Pironti, Pierre-Yves Strub, and Jean Karim Zinzindohoue. A messy state of the union: Taming the composite state machines of TLS. In *2015 IEEE Symposium on Security and Privacy*, pages 535–552. IEEE Computer Society Press, May 2015.
- [14] Karthikeyan Bhargavan, Antoine Delignat-Lavaud, Cédric Fournet, Alfredo Pironti, and Pierre-Yves Strub. Triple handshakes and cookie cutters: Breaking and fixing authentication over TLS. In *2014 IEEE Symposium on Security and Privacy*, pages 98–113. IEEE Computer Society Press, May 2014.
- [15] Karthikeyan Bhargavan and Gaëtan Leurent. On the practical (in-)security of 64-bit block ciphers: Collision attacks on HTTP over TLS and OpenVPN. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *ACM CCS 2016*, pages 456–467. ACM Press, October 2016.
- [16] Karthikeyan Bhargavan and Gaëtan Leurent. Transcript collision attacks: Breaking authentication in TLS, IKE and SSH. In *NDSS 2016*. The Internet Society, February 2016.
- [17] Bitwarden. Bitwarden: About Us. <https://bitwarden.com/about/>. Accessed: 2025-08-26.
- [18] Bitwarden. How End-to-End Encryption Paves the Way for Zero Knowledge - White Paper. <https://bitwarden.com/pdf/resources-zero-knowledge-encryption-white-paper.pdf>, 2025. Accessed: 2025-03-25.
- [19] Lara Bruseghini, Daniel Huigens, and Kenneth G. Paterson. Victory by KO: Attacking OpenPGP using key overwriting. In Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi, editors, *ACM CCS 2022*, pages 411–423. ACM Press, November 2022.
- [20] Michael Carr and Siamak F. Shahandashti. Revisiting security vulnerabilities in commercial password managers. In Marko Hölbl, Kai Rannenberg, and Tatjana Welzer, editors, *ICT Systems Security and Privacy Protection - 35th IFIP TC 11 International Conference, SEC 2020, Maribor, Slovenia, September 21-23, 2020, Proceedings*, volume 580 of *IFIP Advances in Information and Communication Technology*, pages 265–279. Springer, 2020.
- [21] Stephen Checkoway, Jacob Maskiewicz, Christina Gorman, Joshua Fried, Shaanan Cohny, Matthew Green, Nadia Heninger, Ralf-Philipp Weinmann, Eric Rescorla, and Hovav Shacham. A systematic analysis of the juniper dual EC incident. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *ACM CCS 2016*, pages 468–479. ACM Press, October 2016.
- [22] Graeme Connell, Vivian Fang, Rolfe Schmidt, Emma Dauterman, and Raluca Ada Popa. Secret key recovery in a Global-Scale End-to-End encryption system. In *18th USENIX Symposium on Operating Systems Design and Implementation (OSDI 24)*, pages 703–719, Santa Clara, CA, July 2024. USENIX Association.
- [23] Dashlane. 2023 Report - Password Health Scores Around the World. <https://www.dashlane.com/uploads/2023/10/2023-Global-Password-Health-Score-Report-1.pdf>. Accessed: 2025-08-26.
- [24] Dashlane. Dashlane’s Security Principles & Architecture. <https://www.dashlane.com/download/whitepaper-en.pdf>. Accessed: 2025-03-25.
- [25] Gareth T. Davies, Sebastian H. Faller, Kai Gellert, Tobias Handirk, Julia Hesse, Máté Horváth, and Tibor Jager. Security analysis of the WhatsApp end-to-end encrypted backup protocol. In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023, Part IV*, volume 14084 of *LNCS*, pages 330–361. Springer, Cham, August 2023.
- [26] Rose de Fremery. How Zero Knowledge Keeps Passwords Safe. LastPass Blog, 2023.
- [27] Whitfield Diffie and Martin E. Hellman. New directions in cryptography. *IEEE Transactions on Information Theory*, 22(6):644–654, 1976.
- [28] Yihe Duan, Ding Wang, and Yanduo Fu. Security analysis of master-password-protected password management protocols. In Marina Blanton, William Enck, and Cristina Nita-Rotaru, editors, *2025 IEEE Symposium on Security and Privacy*, pages 701–719. IEEE Computer Society Press, May 2025.
- [29] Thai Duong and Juliano Rizzo. Cryptography in the web: The case of cryptographic design flaws in asp.net. In *2011 IEEE Symposium on Security and Privacy*, pages 481–489. IEEE Computer Society Press, May 2011.
- [30] Andrés Fábrega, Armin Namavari, Rachit Agarwal, Ben Nassi, and Thomas Ristenpart. Exploiting leakage in password managers via injection attacks. In Davide Balzarotti and Wenyuan Xu, editors, *USENIX Security 2024*. USENIX Association, August 2024.
- [31] Sascha Fahl, Marian Harbach, Marten Oltrogge, Thomas Muders, and Matthew Smith. Hey, you, get off of my clipboard - on how usability trumps security in android password managers. In Ahmad-Reza Sadeghi,

- editor, *FC 2013*, volume 7859 of *LNCS*, pages 144–161. Springer, Berlin, Heidelberg, April 2013.
- [32] Christina Garman, Kenneth G. Paterson, and Thyla van der Merwe. Attacks only get better: Password recovery attacks against RC4 in TLS. In Jaeyeon Jung and Thorsten Holz, editors, *USENIX Security 2015*, pages 113–128. USENIX Association, August 2015.
- [33] Paolo Gasti and Kasper Bonne Rasmussen. On the security of password manager database formats. In Sara Foresti, Moti Yung, and Fabio Martinelli, editors, *ESORICS 2012*, volume 7459 of *LNCS*, pages 770–787. Springer, Berlin, Heidelberg, September 2012.
- [34] Jessica Gentles, Mason Fields, and Garrett Goodman and Suman Bhunia. Breaking the Vault: A Case Study of the 2022 LastPass Data Breach. <http://www.arxiv.org/pdf/2502.04287>. Accessed: 2025-03-25.
- [35] Nadia Heninger and Keegan Ryan. The hidden number problem with small unknown multipliers: Cryptanalyzing MEGA in six queries and other applications. In Alexandra Boldyreva and Vladimir Kolesnikov, editors, *PKC 2023, Part I*, volume 13940 of *LNCS*, pages 147–176. Springer, Cham, May 2023.
- [36] Jonas Hofmann and Kien Tuong Truong. End-to-end encrypted cloud storage in the wild: A broken ecosystem. In *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security, CCS '24*, page 3988–4001, New York, NY, USA, 2024. Association for Computing Machinery.
- [37] Jonas Hofmann and Kien Tuong Truong. End-to-end encrypted cloud storage in the wild: A broken ecosystem. In Bo Luo, Xiaojing Liao, Jun Xu, Engin Kirda, and David Lie, editors, *ACM CCS 2024*, pages 3988–4001. ACM Press, October 2024.
- [38] Tibor Jager, Kenneth G. Paterson, and Juraj Somorovsky. One bad apple: Backwards compatibility attacks on state-of-the-art cryptography. In *NDSS 2013*. The Internet Society, February 2013.
- [39] Brian Krebs. Password manager LastPass warns of breach. *Krebs on Security*, June 2015. Accessed: 2025-08-25.
- [40] LastPass. LastPass Business. <https://www.lastpass.com/products/team-password-manager>. Accessed: 2025-03-25.
- [41] LastPass. LastPass Whitepaper. https://support.lastpass.com/s/document-item?language=en_US&bundleId=lastpass&topicId=LastPass/lastpass_technical_whitepaper.html&LANG=enus. Accessed: 2025-03-25.
- [42] LastPass. Knowledge workers have a false sense of password security. <https://www.lastpass.com/-/media/7e2b8b715dc44d23ae21c5cf94dc4876.pdf>, 2025. LastPass Business Report; Accessed: August 25, 2025.
- [43] Sean Lawlor and Kevin Lewi. Deploying key transparency at WhatsApp. <https://engineering.fb.com/2023/04/13/security/whatsapp-key-transparency/>, 2023.
- [44] Zhiwei Li, Warren He, Devdatta Akhawe, and Dawn Song. The emperor’s new password manager: Security analysis of web-based password managers. In Kevin Fu and Jaeyeon Jung, editors, *USENIX Security 2014*, pages 465–479. USENIX Association, August 2014.
- [45] NordPass. People have around 170 passwords on average, study shows. <https://www.globenewswire.com/news-release/2024/05/21/2885556/0/en/People-have-around-170-passwords-on-average-study-shows.html>, May 2024. Press Release, Commissioned by NordPass, Conducted by Cint. Accessed: 2025-08-25.
- [46] Sean Oesch and Scott Ruoti. That was then, this is now: A security evaluation of password generation, storage, and autofill in browser-based password managers. In Srdjan Capkun and Franziska Roesner, editors, *USENIX Security 2020*, pages 2165–2182. USENIX Association, August 2020.
- [47] Matteo Scarlata, Matilda Backendal, and Miro Haller. MFKDF: Multiple factors knocked down flat. In Davide Balzarotti and Wenyuan Xu, editors, *USENIX Security 2024*. USENIX Association, August 2024.
- [48] Matteo Scarlata, Giovanni Torrisi, Matilda Backendal, and Kenneth G. Paterson. Zero knowledge (about) encryption: A comparative security analysis of three cloud-based password managers. *Cryptology ePrint Archive*, <https://eprint.iacr.org/2026/58>, 2026.
- [49] Security.org Team. Password Manager Annual Report. <https://www.security.org/digital-safety/password-manager-annual-report/>, 2024. Also titled: Password Manager Industry Report and Market Outlook (2023-2024). Accessed: 2024-11-14.
- [50] David Silver, Suman Jana, Dan Boneh, Eric Yawei Chen, and Collin Jackson. Password managers: Attacks and defenses. In Kevin Fu and Jaeyeon Jung, editors, *USENIX Security 2014*, pages 449–464. USENIX Association, August 2014.
- [51] Ben Stock and Martin Johns. Protecting users against XSS-based password manager abuse. In Shiho Moriai,

Trent Jaeger, and Kouichi Sakurai, editors, *ASIACCS 14*, pages 183–194. ACM Press, June 2014.

- [52] Karim Toubba. Security incident update and recommended actions, 2023.
- [53] Mathy Vanhoef and Frank Piessens. All your biases belong to us: Breaking RC4 in WPA-TKIP and TLS. In Jaeyeon Jung and Thorsten Holz, editors, *USENIX Security 2015*, pages 97–112. USENIX Association, August 2015.
- [54] Serge Vaudenay. Security flaws induced by CBC padding - applications to SSL, IPSEC, WTLS... In Lars R. Knudsen, editor, *EUROCRYPT 2002*, volume 2332 of *LNCS*, pages 534–546. Springer, Berlin, Heidelberg, April / May 2002.
- [55] Rui Zhao, Chuan Yue, and Kun Sun. A security analysis of two commercial browser and cloud based password managers. In *International Conference on Social Computing, SocialCom 2013, SocialCom/PASSAT/Big-Data/EconCom/BioMedCom 2013, Washington, DC, USA, 8-14 September, 2013*, pages 448–453. IEEE Computer Society, 2013.
- [56] Yuliang Zheng. Digital signcryption or how to achieve $\text{cost}(\text{signature} \ \& \ \text{encryption}) \ll \text{cost}(\text{signature}) + \text{cost}(\text{encryption})$. In Burton S. Kaliski, Jr., editor, *CRYPTO'97*, volume 1294 of *LNCS*, pages 165–179. Springer, Berlin, Heidelberg, August 1997.

C Justification of the Threat Model

In addition to the statements made by the password manager vendors about their security goals and guarantees, we give here three further arguments to justify the analysis of cloud-based password managers in the malicious server threat model.

PASSWORDS ARE HIGH VALUE DATA. Passwords are among the most sensitive kind of user data, as they often grant access to accounts and services which may in turn reveal more data, or allow an attacker to impersonate the victim. Since password manager servers host high concentrations of passwords from thousands of users, they thus provide an attractive target for hackers. Attacks on the provider server infrastructure can be prevented by carefully designed operational security measures, but it is well within the bounds of reason to assume that these services are targeted by sophisticated nation-state-level adversaries, for example via software supply-chain attacks [21] or spear-phishing. Moreover, some of the service providers have a history of being breached – for example LastPass suffered breaches in 2015 and 2022, and another serious security incident in 2021. In the 2022 breach, the attacker gained access to the development environment, source code,

and technical information, used that information to compromise the computer of a senior DevOps engineer, then gained access to an encrypted corporate vault, and finally to the keys of the Amazon S3 “buckets” of the backups to customer files [34, 52]. In a self-hosted deployment involving a less technically advanced customer, it seems highly likely that the server hosting user vaults could not resist such a determined attack. While none of the breaches we are aware of involved reprogramming the server to make it undertake malicious actions, this goes just one step beyond attacks on password manager service providers that have been documented. Active attacks more broadly have been documented in the wild.¹⁰

COMPARISON TO OTHER E2EE DOMAINS. Security in the face of a malicious server is by now a standard expectation in related settings such as messaging systems and cloud storage, where sensitive user data (in the form of private messages or files) is also relayed or stored by service providers. In particular, the malicious server threat model is the accepted interpretation of the security guarantees provided by end-to-end encryption in these settings. (We note that vendors of E2EE cloud storage also use “Zero Knowledge Encryption” to refer to E2EE. Thus, also in this setting, this term is taken to mean data confidentiality and integrity in the face of a malicious or compromised server.) For example, recent analyses of cloud storage systems [3, 4, 11, 35, 37] have forced vendors in this space to upgrade their approach to provide security against a malicious server, while major messaging systems like Signal, WhatsApp and iMessage have long featured E2EE in the standard sense (Telegram remains a notable and regrettable exception in this area, with E2EE only being optionally available on a per conversation basis rather than as a default, and not available at all for groups). We think it is reasonable to expect these terms (zero knowledge and E2E encryption) to have the same meaning in the context of password managers.

QUANTIFYING THE PROVIDED SECURITY. Given the sensitivity of the data they store and the scale at which they operate, we need to understand the security limits of password managers. This requires us to answer the question of which threats they can withstand, and which they cannot. Through this, we are able to identify areas where their practical security can still be improved. This calls for analysis in a range of threat models, including a strong threat model like ours. Eventually, we expect the community to develop appropriate formal models and security proofs for password managers, similar to those for cloud storage presented in [10].

In summary, we consider the malicious server setting to be an appropriate threat model in which to analyse the security of password managers.

¹⁰ See, for example, Quantum Insert <https://theintercept.com/2014/03/12/nsa-plans-infect-millions-computers-malware/>) and the more recent Secret Blizzard <https://www.microsoft.com/en-us/security/blog/2025/07/31/frozen-in-transit-secret-blizzards-aitm-campaign-against-diplomats/>.