LABRYS FRONTIER SERIES

Password management:

Where LastPass got it wrong and how to approach the issue in organizations

March
2023

https://doi.org/10.59262/ce3fc7

Password management:

Where LastPass got it wrong and how to approach the issue in organizations

Labrys

Jonas de Abreu
Mariana Cunha e Melo

Labrys's mission is to connect technology, business, and policy to build, inspire, and enable society-focused technology at scale. Our ultimate goal is to create a world where technologists can develop society-focused technology, citizens can effectively debate what society needs from technology, and regulators can align incentives to strengthen an open, inclusive, and sustainable economy for all.

Center for Technology and Public Interest, SL
Carrer de Bailèn, 11, Barcelona, Spain, 08010
www.wearelabrys.com

Table of contents

Summary

In August 2022, LastPass suffered two cyberattacks that breached customer data and encrypted passwords. LastPass acknowledged the attacks, but their communication was not transparent enough. In November 2022, a follow-up attack compromised customer data further. LastPass communicated that this was a low-risk attack and that customers did not need to take any action. However, in December 2022, LastPass admitted the actual scale of the breach, and that all customer vaults were compromised. It is important to note that every company suffers frequent attacks, but the proper security posture under this type of attack is to assume that everything will eventually get compromised. The incident makes a case for why companies should always deploy additional defenses, such as employing security keys, to stay secure in the long term.

I. The LastPass incident

Starting in August 2022, LastPass suffered two successful cyber-attacks that resulted in a massive breach of customer data and encrypted passwords.

Before we get into our analysis and recommendations, we should highlight something. The occurrence of cyber-attacks (whether successful or not) is not a good enough reason to switch your business to a different company. We can't stress it enough: every - and we mean every - single company suffers frequent attacks - even if they don't know or acknowledge it. Our point here is that the way LastPass handled and communicated these attacks to their customers raised several flags for us. It is important to disclose that we were old-time users of their services and,

after finishing this analysis, we decided to move to a different solution.

We'll detail the rationale for this decision, as there are many helpful learning opportunities on how not to do things in case of a security incident.

I.1. The facts

On August 25th last year, LastPass said:

> "We have determined that an unauthorised party gained access to portions of the LastPass development environment through a single compromised developer account and took portions of source code and some proprietary LastPass technical information. Our products and services are operating normally".

That communication made us trust LastPass more, not less. Public acknowledgment of a breach should be received as something that speaks highly of your provider, especially for security-sensitive companies. Every company gets breached, and the more they are transparent about it, the clearer the message that they take security seriously is. It is not by chance that GDPR mandate that personal data breaches are notified, and countries have vulnerability disclosure programs, such as the one run by CISA.

So when LastPass said:

> "**4. What should I do to protect myself and my vault data?** At this time, we don't recommend any action on behalf of our users or administrators. As always, we recommend

> that you follow our best practices around setup and configuration of LastPass which can be found here."

We had no reason to be alarmed.

LastPass also clarified that this was a low-risk attack and that customers did not need to take any action. As far as a successful attack goes, this one seemed pretty harmless.

Unfortunately, a more severe follow-up attack changed this scenario, as LastPass disclosed on November 30th:

> "We have determined that an unauthorized party, using information obtained in the August 2022 incident, was able to gain access to certain elements of our customers' information. **Our customers' passwords remain safely encrypted due to LastPass's Zero Knowledge architecture**".

At this point, it was unclear whether customer passwords had been compromised. Still, LastPass addressed this issue directly and affirmed all passwords were safe and sound because of LastPass's Zero Knowledge architecture. That affirmation was strange since there was no mention of any customer's vault (how LastPass calls their encrypted password database) data being accessed.

The degree of seriousness of the November 30th attack only surfaced on December 22nd, when they informed that their encrypted vault data was compromised.

> "The threat actor was also able to copy a **backup of customer vault data** from the encrypted storage container which is stored in

> a proprietary binary format that contains both unencrypted data, such as **website URLs, as well as fully-encrypted sensitive fields such as website usernames and passwords, secure notes, and form-filled data**. [...]
>
> However, it is important to note that if you are a Business customer who is not using Federated Login and your master password does not make use of the defaults above, then it would significantly reduce the number of attempts needed to guess it correctly. In this case, as an extra security measure, you should consider minimising risk by changing passwords of websites you have stored. [...]
>
> We have already **notified a small subset (less than 3%)** of our Business customers to recommend that they take certain actions based on their specific account configurations. If you are a Business customer and you have not already been contacted to take action, then there are no other recommended actions for you to take at this time".

That disclosure substantially changed the threat profile of the November 30th attack. Before the breach, any party attempting to attack a customer needed to go through LastPass' online services before trying to decrypt any vault. That is an important protection layer because it allows LastPass to deploy many protection measures, such as rate-limiting password guess attempts. It was also hard to determine which customers would be the most valuable to attack (large bank accounts, etc.) since the attacker would need to infer the target value only from the email used as a username.

We call this an **Online Attack**. It is one of the best scenarios to defend a system, especially because you can update and improve your defense as attack techniques evolve.

After the breach, the attack scenario changed because now the attacker doesn’t need to go through LastPass services to attempt accessing the encrypted data. So they can make as many attempts and deploy all sophisticated techniques that Laspass' service defenses would otherwise block. It is also impossible to improve defenses because a copy of the data is already in the attacker's hands, and LastPass or the user can't change the properties of the vault anymore. The only layer of protection now preventing the attacker from accessing the data is the cryptography employed at the time and the strength of the customer's master password. And even though the cryptographic algorithms are sound today, that may change tomorrow with better attack techniques, and, again, there is no way of updating the vault.

We call this an **Offline Attack**. The barriers companies put in place to prevent this type of attack are more of a last stand against an attacker than an effective defensive strategy. The proper security posture under this type of attack is to assume that everything will eventually get compromised. You need to deploy additional defenses to stay secure in the long term.

In the December 22nd communication, LastPass admitted that business customers were affected but notified with recommendations only 3% of accounts they considered to

be at a higher risk of compromise. It's worth noting, though, that it is hard to know precisely how many of the business customers had vault data leaked and how many were using passwords that could be compromised. LastPass cannot know it since they never access the customer master password (their Zero-Knowledge architecture). Therefore, deciding to notify _only some of _their business customers that weren't integrated by a Single Sing On solution at the time was odd, at best.

Only three months later, on March 1st, LastPass finally admitted the actual scale of the breach:

> "Cloud-based backup storage - contained configuration data, API secrets, third-party integration secrets, customer metadata, **and backups of all customer vault data**. All sensitive customer vault data, other than URLs, file paths to installed LastPass Windows or macOS software, and certain use cases involving email addresses, were encrypted using our Zero knowledge model and can only be decrypted with a unique encryption key derived from each user's master password. As a reminder, end user master passwords are **never** known to LastPass and are not stored or maintained by LastPass - therefore, they were not included in the exfiltrated data".

The attackers accessed all customer vaults. It is essential to highlight that this was not a third attack. The second attack compromised all backups. Oddly, their monitoring infrastructure pinpointed 3% of their business customers were compromised in December but could not detect the copy of all vaults backups.

They also provided a specific security bulletin for Premium customers:

> "Do you need to act? Ask yourself these questions to decide:
>
> Is your master password strong and unique?
>
> yes / no / unsure
>
> Is your master password hash Iteration value set to at least 600,000?
>
> yes / no / unsure
>
> Are the passwords in your vault all strong and unique?
>
> yes / no / unsure
>
> Are you using multifactor authentication on LastPass and other important accounts?
>
> yes / no / unsure
>
> Did you answer no or unsure to any of these questions? If so, keep reading and please take the recommended actions until all answers are a yes".

None of the mitigations proposed for a **no** or **unsure** answer involves resetting stored passwords. Reading the bulletin, a customer may think that changing the master password can help, but that will only secure them against a future attack. The attack has become an **offline attack** now. The only things that matter to the security of the breached vaults are time and the effort and resources the

attacker will employ to get any given customer's credentials.

This offline attack is dangerous because the attacker possesses all login URLs with passwords in the vaults. The attacker can prioritize attacks against vaults with bank passwords, online shops, phone companies (to bypass SMS OTPs), or whatever else is their fraud expertise. It's like if a burglar obtained the ability to try millions of keys per second on any lock and then were given the address to your home.

Looking at the whole incident and the disclosure approach they took, we decided to stop using LastPass, not because of the incident per se, but because of how they handled it. We lost the confidence that, in a future incident, LastPass would promptly inform us and suggest proper mitigations.

Three things are critical in disclosing a breach: it must be timely, as complete as possible, and provide reasonable mitigations in case the customer can take action. LastPass failed to deliver on all three.

We also had an issue with the writing of the notifications. There was a lot of emphasis on their master password best practices, pointing out that there was a risk only if customers didn't follow them. While that is technically true today, it sends the message that LastPass is trying to shift the blame for the customer being at risk to the customer itself. That is not what they sell on their front page:

> "Password Management from Anywhere

> Life is happening online. Work. Play. Family and friends. LastPass puts your digital life at your fingertips, simply and securely".

Companies should not blame their customers for their failures.

I.2. So, should we stop using password managers?

No. Even though password managers have limitations and risks, as Stuart Schechter and Terence Eden brilliantly show, they are the best solution today for dealing with the complexity of login to most websites.

That said, a password manager cannot secure you by itself.

As mentioned, we were LastPass customers, and the attackers copied our vault data in this breach. But we could calmly migrate to a different provider instead of hushing to sites to change our passwords, even though our position now is to consider all passwords we had at LastPass to be compromised.

The reason is that we were also relying on something other than LastPass. None of our second authentication factors (2FA) resided there, and we enabled it on every site we know supports it. We store those authentication factors in a physical security token (a Yubikey), which we always carry, and on a backup that stays home. They were enough to buy us time to study the current threats we're under, the current state of password managers, and then move to the new one. The security token's additional protection would

be the same, even if we never learned about LastPass' breach.

II. Recommendations for business password management

II.1. Key learnings

There are many things we can learn from this attack. Our two main ones are how devastating this attack can be and how much additional security a physical security key can add, even in the worst-case scenario.

Our most important recommendation is to ensure everyone in your company activates 2FA (in FIDO mode) using a hardware key such as Yubikey or Titan Security Key. If that is not possible, enabling TOTP (those six digits that change every 30 seconds) based on a hardware key is better than with a software application, but that provides fewer protections against attacks.

Phishing and its variants were the most internet-related crime reported to the FBI in 2021 (page 22), and FIDO makes the best current protection against this type of attack.

There is a reason for this solution to be so effective. The FIDO protocol prevents users from making mistakes by relying on cryptographic guarantees. It is easy to fool a human but quite challenging to fool cryptography.

Even though a hardware 2FA is not a substitute for good password management practices, it can provide the best security for the implementation effort because it is an added tool that works well independently of other password management practices.

A note about cost: the price of physical security keys starts at EUR 20.00 and can go over EUR 50.00 depending on the model and manufacturer. Considering a whole company deployment, it can amount to a substantial cost. We argue that when properly used, it is worth the price.

II.2. How CTPI manages passwords

The short answer is that we use Bitwarden as a password manager and Yubikeys for 2FA. Each user generates their Bitwarden master password using EFF's Diceware word list and has 2 Yubikeys: a main one and a backup one.

The longer answer is a bit more complex, and we'll explore the multiple steps of this process, from threat modeling to solution design.

Defining how CTPI was going to manage passwords was one of the things we wanted to have ready before CTPI went live, and now we're going to explore how we did. We focused on making it as easy to use as possible (to prevent us from not wanting to use it) even if we had to sacrifice some security. If it is too hard or bothersome to use daily, the best security possible will not work for us. We don't want to drive non-compliance because of friction.

a. Threat modeling

We started by defining attack scenarios and making all of our assumptions explicit. We would clearly use a password manager and a hardware key since that is the minimum requirement for securing access to multiple web applications.

We found the following that could range from annoyance to catastrophic:

1. Compromise of the password manager (LastPass case)
2. Compromise of the master password
3. Site compromise that results in our passwords leak (Have I Been Pwned list of compromised sites)
4. Targeted attacks against our team members (Spearhead phishing)
5. Non-targeted phishing attacks (UK NCSC phishing prevention guidance)

6. Credential stuffing against sites, based on leaked passwords (OWASP credential stuffing)

7. Physically stealing user-controlled security devices (cellphones, yubikeys, etc.)

We also established under which assumptions our solution should be strong enough. Setting the assumptions is essential because a solution that can face state-sponsored attacks could be substantially different than what we need to protect our team members.

The main assumptions were:

A. We are not targets for state-sponsored attacks

B. If Yubikeys can be compromised, only states know how to do that

C. Physical attacks by skilled adversaries are unlikely

D. Most assets don't need to resist highly skilled targeted attackers, but it would be nice to have

E. The attacker has not compromised our machines (computer, tablet, cellphone). This assumption is about the attacker's skill rather than the trustworthiness of our machine hardening.

b. Solution design

A big part of the solution design process was ensuring we had mitigations for the attack scenarios we were considering.

To mitigate password leaks and credential stuffing **(items 3 and 6),** all the passwords must be unique and hard to guess. We do this by generating different random passwords for each site. Since it is unfeasible to memorize hundreds of random passwords and writing them down has prohibitive usability and security issues, a password manager is necessary.

Considering a possible compromise of the password manager **(item 1),** passwords must not be accessible to the server (they need to be end-to-end encrypted), requiring the user to memorize a master password.

The phishing attack variants **(items 4 and 5)** can be devastating, but fortunately, the FIDO protocol can mitigate it properly. FIDO usually requires specific cryptography hardware, which we can get as a small physical security key.

We need defenses to handle a complete compromise of the password manager. Going with an offline password manager is considered the best security, but the usability implications make them unviable. To control the added risk of an online password manager, we'll have all forms of 2FA handled by the physical token. We should avoid relying on TOTP whenever possible and use FIDO instead. Using this same physical token to control access to the password manager also protects the user against a master password leak since it limits access to the encrypted passwords database in the case of an Online Attack. Usually, this will not improve security against an Offline Attack.

Considering the assumptions about state-level and physical attackers **(items A, B, and C)** and the above requirement of memorizing a master password, a physical security key is acceptable. Carrying a primary security key in person (our suggestion is their home keychain) and having a backup copy hidden in a relatively safe place (our suggestion is their houses) is enough redundancy.

Finally, since having the key stolen **(item 7)** is unlikely to be part of a cyber attack, we decided we didn't need to focus on specific mitigation. Using the backup key to reset 2FA on all sites should be enough if the user loses control of the security key.

Considering all this, this is our solution:

- CTPI provides a Bitwarden account and two Yubikeys to all team members. When creating the Bitwarden account, the user generates the master password based on the EFF's Diceware word list.
- The team member writes this password on a small piece of paper and keeps it in their wallet until it is memorized. After memorization, it should be destroyed.
- To add an account, the user uses Bitwarden to generate a secure random password for that account and then uses both Yubikeys to enable 2FA for that account.

We are a small team now, so we have multiple people managing Yubikey backups to ensure enough redundancy.

For more extensive deployments where various people have admin access to services, it's enough to have a single Yubikey per person and rely on 2FA reset by other admins if one loses control of their security key.

About the authors

Jonas de Abreu is an experienced security engineer with a strong background in software engineering and security. He worked as an external consultant and instructor in the early stages of his career. Jonas spent almost a decade at Nubank, the leading fintech company in LATAM, where he made significant contributions as the Chief Information Security Officer and later as a Principal Engineer. He played a crucial role in shaping Nubank's infosec team and strategy and made a notable impact on the Brazilian Central Bank's decisions regarding the Pix rail. Jonas was responsible for Nubank's technical proposals during the Pix Forum, which were highly valued by industry experts and had a pivotal role in shaping the future of payment systems in Brazil.

About the authors

Mariana Cunha e Melo is an experienced lawyer with a background in strategic litigation, public policy, legal research, and regulated markets. She has worked with Supreme Court justices in Brazil, represented Google in high courts and strategic litigation cases, and built the internet law team at a top law firm. At Nubank, she helped structure the Public Policy team and led efforts to work with the Brazilian Central Bank on designing its Real-Time Payments rail. She has also worked in early-stage startups as a strategic projects owner and director of regulation and strategy. Mariana is also a writer and speaker on topics such as privacy, free speech online, and regulation of fintech companies, with numerous international events under her belt.