Unveiling the Shadows:

Understanding Token Theft

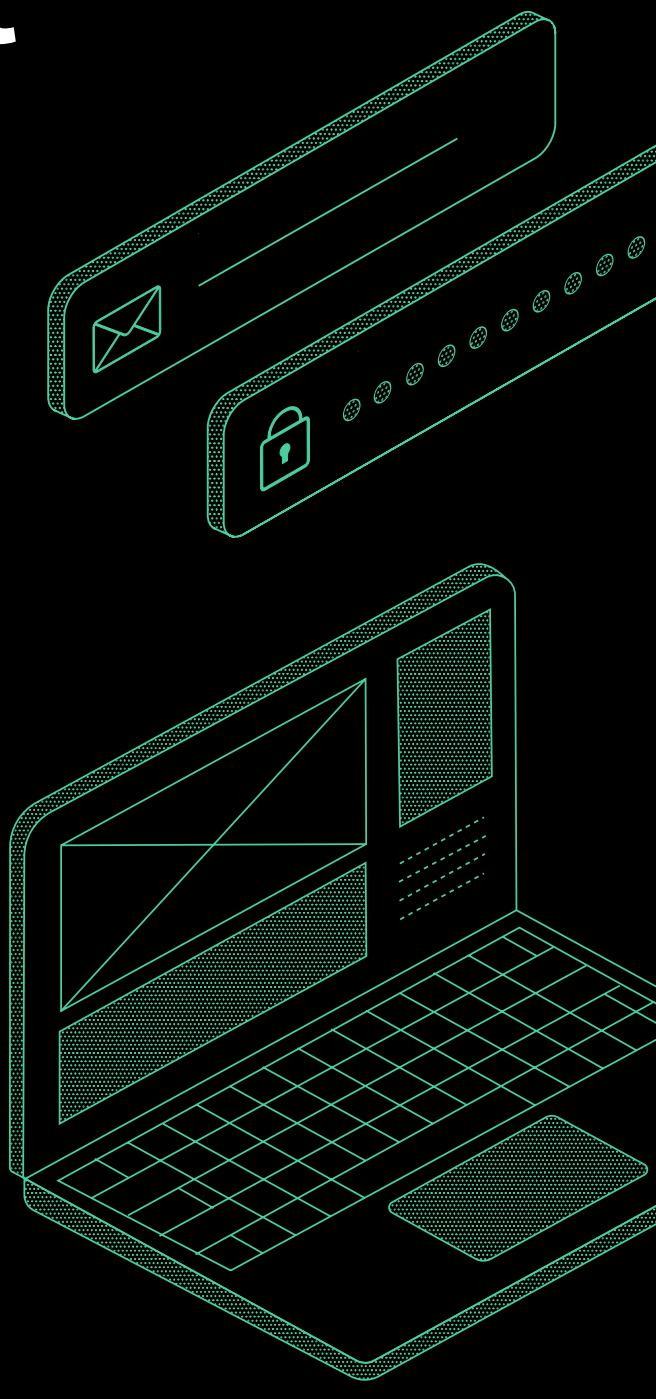

Triskele Labs

Authors

Brecht Snijders
Principal Offensive Consultant
Operations Belgium

Jason Trapp
DFIR Analyst
Operations Canada

Caleb Boyd
DFIR Analyst
Operations Australia

Cameron Paddy
DFIR Analyst
Operations NZ

Mike Varley
DFIR Analyst
Operations UK

In this whitepaper, the Triskele Labs Digital Forensics and Incident Response (DFIR) and Offensive Teams delve into the intricacies of token theft, the motivations behind the threat groups doing it, and scenarios from real life DFIR investigations showcasing the most prevalent methods. To wrap up, we will explore the proactive measures that organisations can take to fortify their defences against this particular threat.

> **The transmission and storage of these tokens must have appropriate measures to protect them to prevent Threat Actors from exploiting them for their own gain.**

Content

Introduction

As the cloud continues becoming more ubiquitous, more organisations are migrating from on-premises solutions to the cloud. A prime example of this would be businesses transitioning from on-premises Microsoft Exchange to M365.

When a user logs into online services (like accessing your work email account from Outlook.com) these services will use a login token which contains a unique identifier to ensure sessions are valid for a set period of time. If you are returning to a web site after logging in previously, and get logged in straight away, without having to authenticate, these are tokens at work.

These login tokens are generated by servers to authenticate and authorise user access to that service. The transmission and storage of these tokens must have appropriate measures to protect them to prevent Threat Actors from exploiting them for their own gain.

Token theft occurs when a Threat Actor gains unauthorised access to login tokens which would result in the Threat Actor being able to impersonate a legitimate user. Where token theft is particularly alarming, is when it can be used to bypass Multi-Factor Authentication (MFA).

In this whitepaper, the Triskele Labs Digital Forensics and Incident Response (DFIR) and Offensive Teams delve into the intricacies of token theft, the motivations behind the threat groups doing it, and scenarios from real life DFIR investigations showcasing the most prevalent methods. To wrap up, we will explore the proactive measures that organisations can take to fortify their defences against this particular threat.

Token theft occurs when a Threat Actor gains unauthorised access to login tokens which would result in the Threat Actor being able to impersonate a legitimate user.

What is Token Theft

The Cliff Notes

In simple terms, token theft occurs when a Threat Actor gains access to an authentication token used by an online service. Authentication tokens are pieces of information stored locally on a device and act as a mechanism to tell an online service that this device has authenticated recently enough that it does not need to do so again.

Token theft can occur through multiple methods. The most common (and the one we demonstrate here) is an adversary-in-the-middle attack where a Threat Actor will create a malicious website which they will direct a victim to, often through a phishing email, where the victim will be asked to log in.

Because the attacker is positioned in the middle of this activity, they're able to see both what you send the site you think you're logging into (your login information), as well as what is being returned by that site (authentication token). If the login information supplied is correct, the legitimate site will return a token that can be used to authenticate future logins. Token theft is when the Threat Actor intercepts and obtains this token.

The most common (and the one we demonstrate here) is an adversary-in-the-middle attack where a Threat Actor will create a malicious website which they will direct a victim to, often through a phishing email, where the victim will be asked to log in.

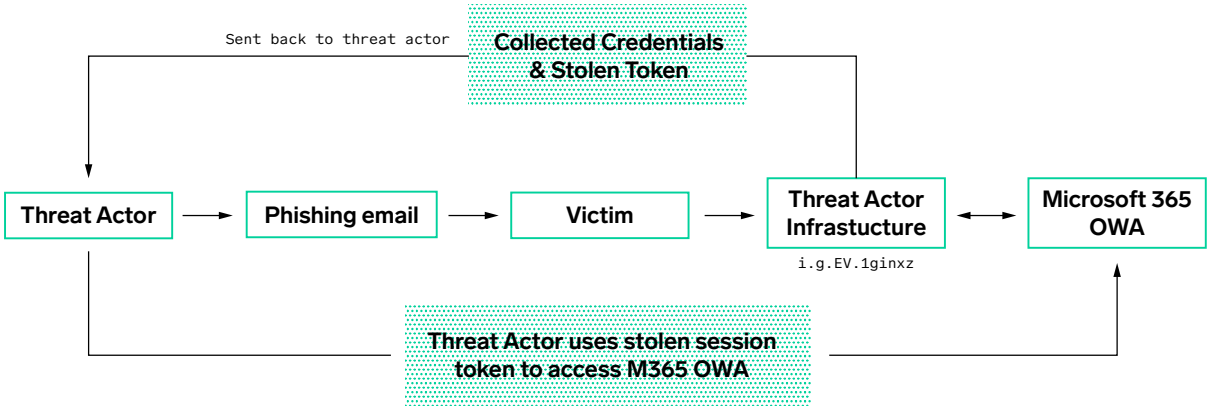

How do Threat Actors Steal Tokens?

After responding to multiple DFIR engagements where tokens had been stolen to access user accounts, the DFIR team decided to work with the Offensive team to stand up some common tooling being used by Threat Actors to achieve these attacks.

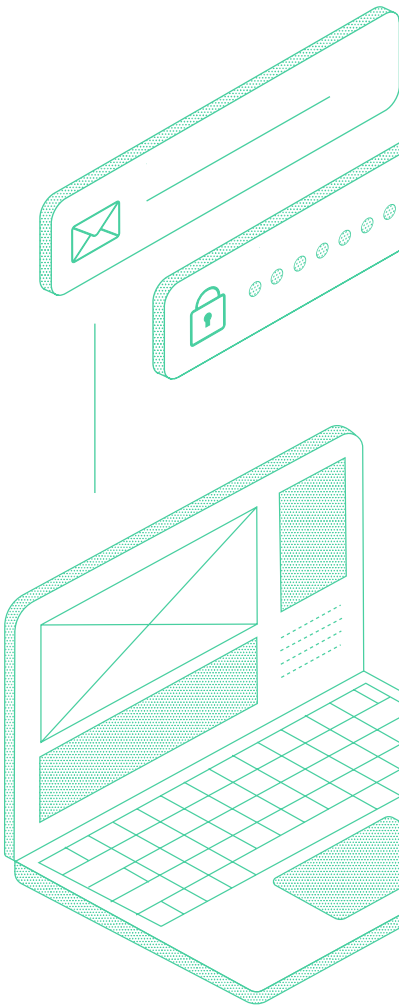

Services that utilise login tokens will differ in their implementation. In this instance, the following platforms and tools were used:

- Microsoft 365 (M365) E3 plan,
- Evilginx2 hosted externally,
- Edit-Cookie (Firefox Extension).

Evilgnix is a tool that can be used by Threat Actors to host their own copies of websites to perform a man in the middle attack. This allows them to capture the information of a victim while still passing it on to the legitimate service, allowing them to capture login details and any returned session token if the authentication is successful.

Edit-Cookie is a Firefox extension that allows you to add your own or edit existing cookies for a website.

Within M365, a new domain and mailbox were created to receive emails. This mailbox was created to simulate what the intended target of phishing scam would see if they were targeted by a Threat Actor.

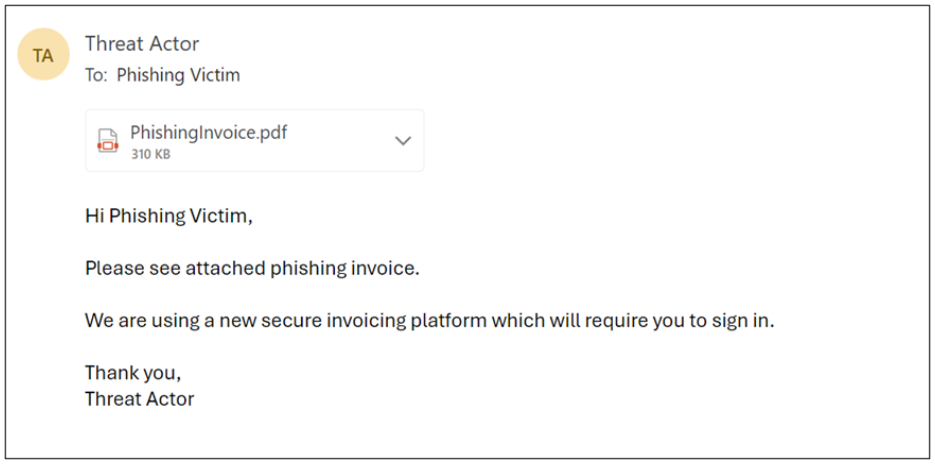

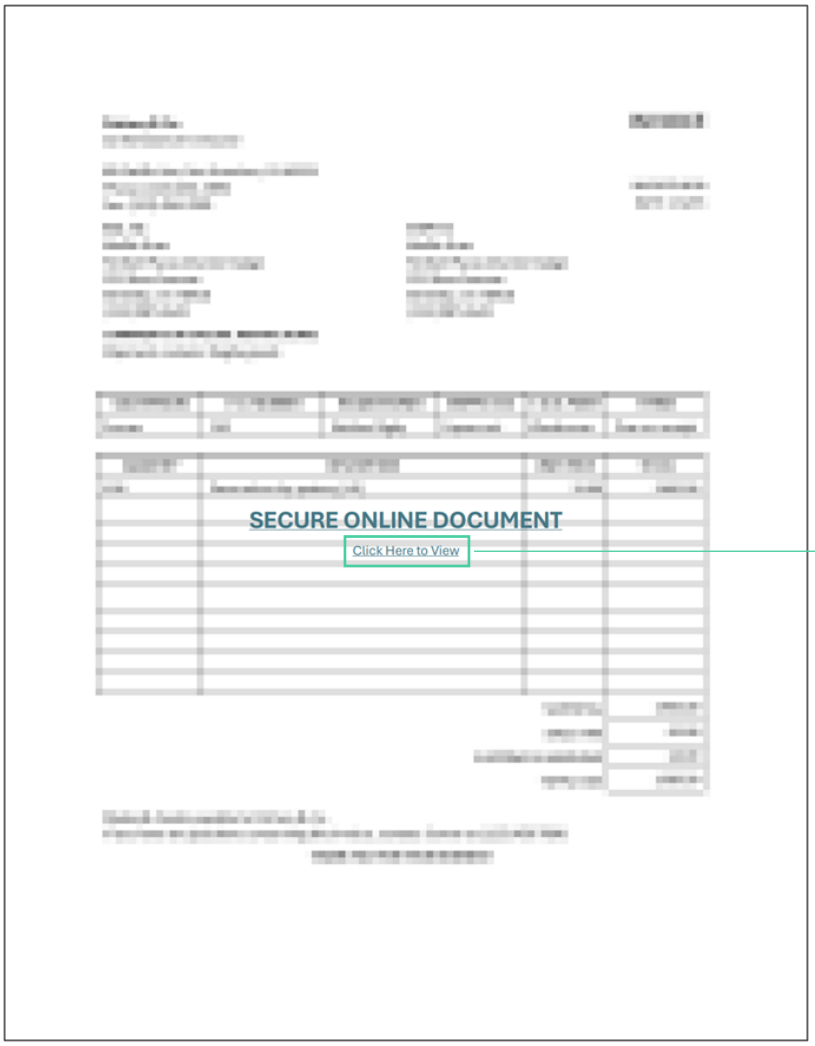

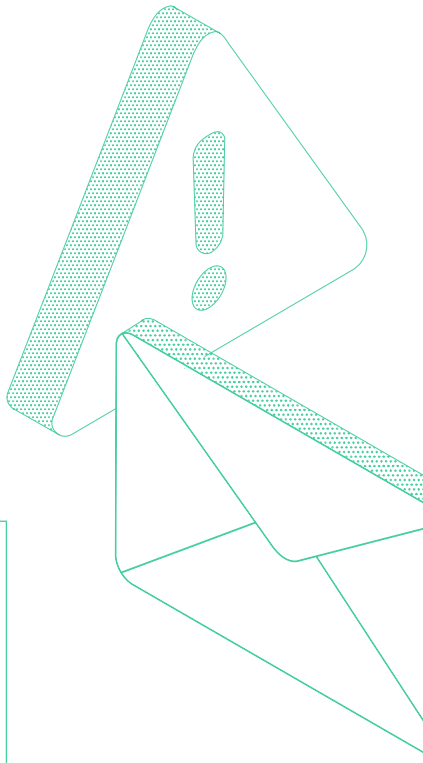

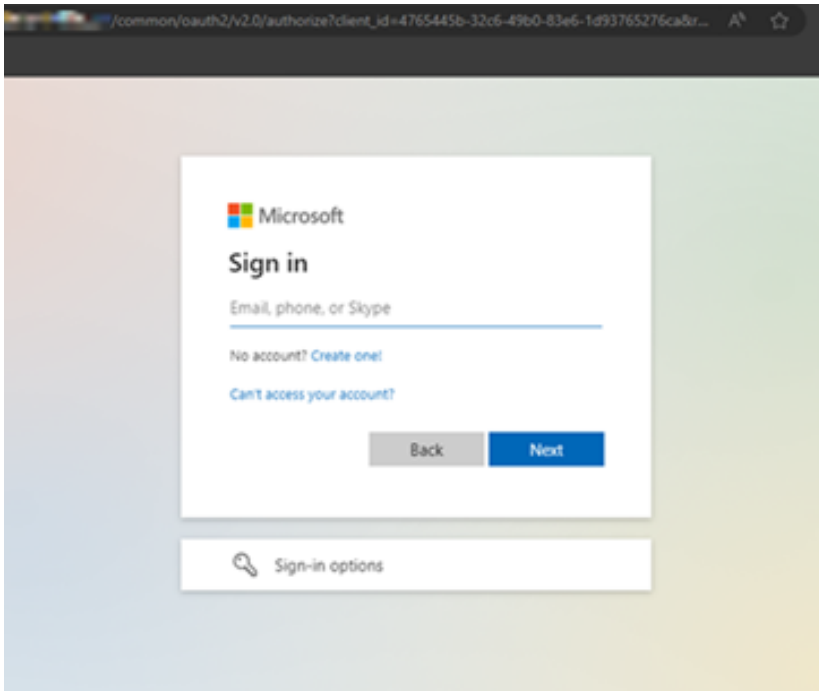

From the above we can see the phishing email contained a PDF which had a link to the phishing site that was setup by the team. When the link was clicked, the theoretical victim was presented with a standard looking Microsoft 365 login page.

Once the user entered their credentials into the phishing page, this information was sent to the Evilginx server which collected the session token.

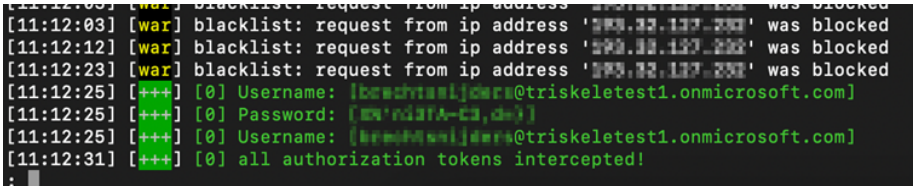

During the first test, the victim account did not have MFA configured — the next step was to apply MFA to the account with the Microsoft Authenticator app.

Did MFA configuration prevent the Threat Actor from stealing enough information to login? Unfortunately, the answer to this question was no, MFA did not make a difference in the ability to steal the session token, as shown below.

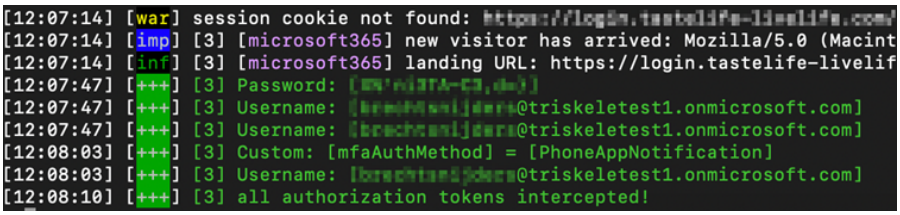

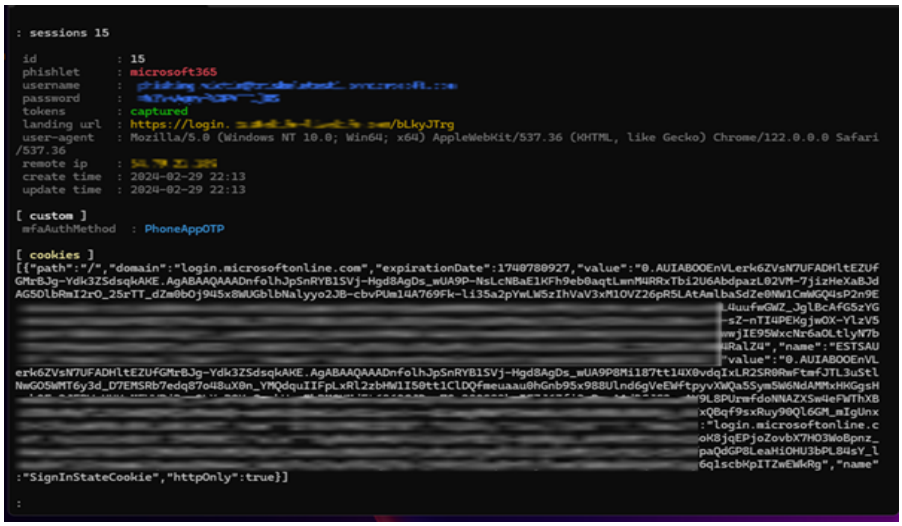

Threat Actors and their tools have grown more sophisticated and the Evilginx server hosting the phishing site essentially ferried the authentication details, including MFA, to the legitimate provider on behalf of the user and, upon successful authentication, stole the session token and redirected the victim to the legitimate service provider.

> Did MFA configuration prevent the Threat Actor from stealing enough information to login?

Now in possession of the session token, the Threat Actor can then use a number of methods to inject this information into their chosen web browser and navigate to the legitimate service providers' webpage. The legitimate service provider will then allow the Threat Actor to bypass authentication because, to the service provider, it appears as though they have pre-authenticated.

The legitimate service provider will then allow the Threat Actor to bypass authentication because, to the service provider, it appears as though they have pre-authenticated.

As a note, the level of security provided by different MFA types are not equal. For instance, SMS, voice, or email one-time passwords MFA methods can be intercepted by a Threat Actor allowing them to gain access to the account. App based MFA tokens such as those provided by the Microsoft Authenticator, Google Authenticator, or others similar apps are much less likely to be intercepted by a Threat Actor and are more secure methods of MFA, but are not immune to session interception, as we can see from the screenshot above where MFA was used but the session was still intercepted by the Threat Actor.

From the screenshot below we can see that it is difficult to identify token theft, with the only indications being three different IP addresses and locations and the User Agents of the browsers all used within short succession. We can see our victim favours the Google Chrome browser, while the Threat Actor is using the Mozilla Firefox browser (for demonstration).

Date (UTC)	Username	Unique Token Identifier	IP address	Location	Status	Browser	Multifactor authentification result
29/02/2024 21:30	phishing.victim@domain.com	YEIMWIXIHkCIja-leHbRpAA	54.X.X.X	Sydney, NSW, AU	Success	Chrome 122.0.0	MFA requirement satisfied by claim in the token
29/02/2024 22:13	phishing.victim@domain.com	6HZ7cYab-d0G1EppHaNuiAA	170.X.X.X	Lincolnshire, Illinois, US	Success	Chrome 122.0.0	MFA requirement satisfied by claim in the token
29/02/2024 22:17	phishing.victim@domain.com	XjAWLG-vDrEu5drlGYvrbAA	103.X.X.X	Petrie Terrace, QLD, AU	Success	Firefox 123.0	MFA requirement satisfied by claim in the token
29/02/2024 22:44	phishing.victim@domain.com	WEkj6a1hsEN-GWEBh4jnAA	54.X.X.X	Sydney, NSW, AU	Success	Chrome 122.0.0	MFA requirement satisfied by claim in the token

Logs:

- 21:30 – Victim logging into their account from their normal working location
- 22:13 – Victim trying to login to the Threat Actors phishing site
- 22:17 – Threat Actor using the stolen token to log into the account
- 22:44 – Victim logging into their account from their normal working location

Outside the IP address being different there is no other indication the account may have been comprised by the Threat Actor. There is no other information available in the Microsoft logs to suggest that the token that was stolen at 22:13 was then used at 22:17 by the Threat Actor.

Identifying a Malicious Website

Threat Actors will use a myriad of techniques to entice a legitimate user to open the website or attachment within their email. Users within your organisation should look for the following before interacting with any email:

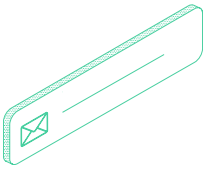

Sender's email address – Threat Actors will often use lookalike email addresses to trick employees that the email is coming from a legitimate source. Due to the spacing between characters of an email address being small, it is common for Threat Actors to create a lookalike domain by using characters that look similar to the original address such as replacing the letter 'm' with 'rn'.

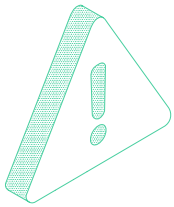

Urgency – Threat Actors will often pressure users into clicking a link as the offer or email verification link is only good for the next 30 minutes, as an example.

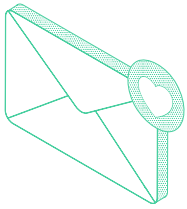

Emails are too good to be true – Some malicious emails will state things that are too good to be true such as gift cards. In one incident, the Threat Actor informed the user they were getting a pay rise.

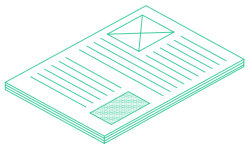

Emails that request you to ignore policies – Most organisations will have policies around payment details being changed, however, Threat Actors will attempt to convince employees to ignore these details to their own.

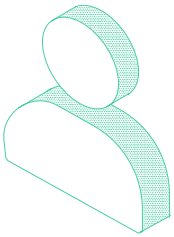

Emails from someone you know - These emails can be detected as the wording or tone may appear different to what you would usually expect.

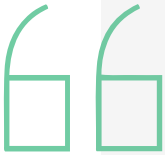

Threat Actors will often pressure users into clicking a link as the offer or email verification link is only good for the next 30 minutes, as an example.

Comparing a Legitimate Website from a Threat Actor Site

Often Threat Actors will create their own website which is a replica (or as close as they can get it) of the legitimate site which they will then send to the person they are wanting to compromise, that person then logs in thinking it's the actual site. From this the Threat Actor is then able to steal the username, password and the authentication token that was generated.

Since these websites are often a clone of legitimate ones, they're often difficult to spot. Below are some screenshots, the first is of the actual Microsoft sign-in page and the second is of the website the Triskele Labs team created to demonstrate Token Theft.

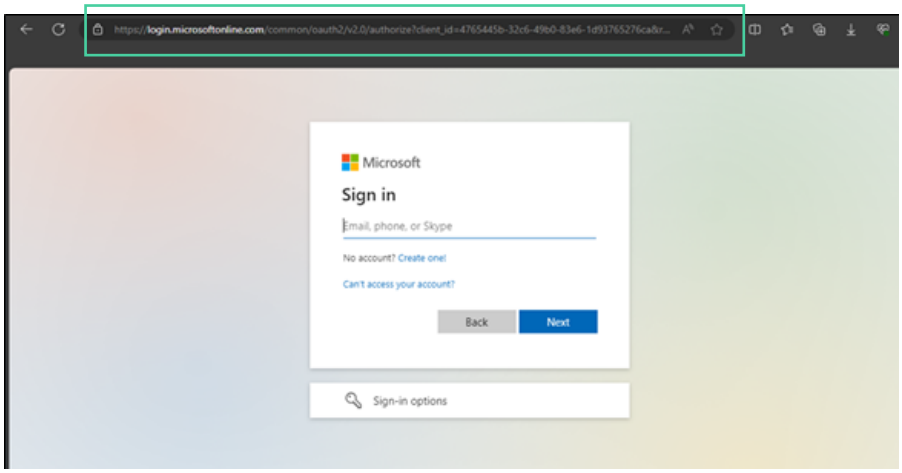

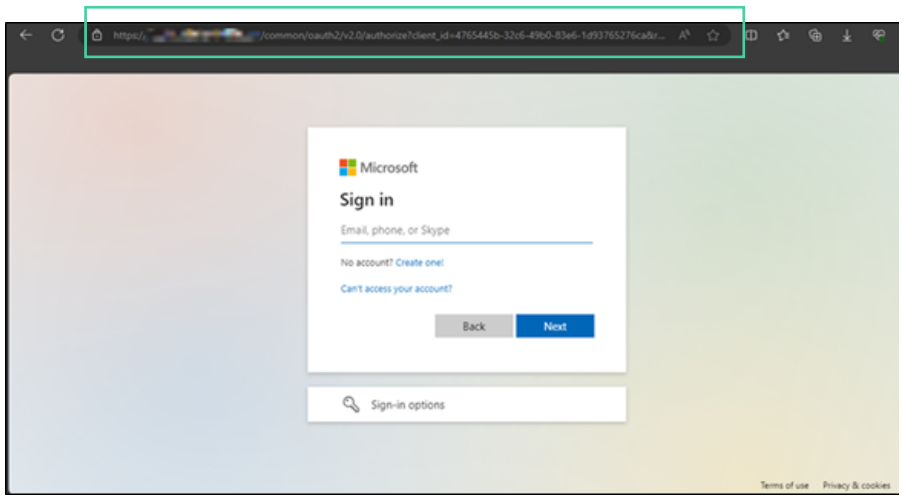

Comparing the two screenshots we can see the pages themselves look identical, and the phishing website did have some odd behaviour once logged in and took longer to respond to requests than normal, but by then it was too late,

> Since these websites are often a clone of legitimate ones, they're often difficult to spot.

the credentials and the token had already been taken by the attacker. The only real difference between the two sites that made it obvious one was a phishing site was the URL, with the phishing website being the Triskele Labs test domain, rather than microsoftonline.com.

It's particularly important to pay close attention to the URL as Threat Actors will often try to make their malicious URL appear as close to the legitimate ones as possible.

> The only real difference between the two sites that made it obvious one was a phishing site was the URL.

Detecting and Remediating Token Theft

Whilst credential theft poses a significant risk to any environment, there are multiple methods that can be used to detect and prevent credential theft within an environment.

Phishing Resistant Devices – FIDO 2.0 compliant devices such as the YubiKey are known for being phishing resistant devices.

Security Awareness Training – By training employees of the dangers that malicious emails pose to the organisation, this decreases the risk of users clicking on malicious emails.

Unified Audit Logging – This feature within M365 is able to capture logs from different services across your tenancy such as, SharePoint, Azure, Exchange, and more. It is important to note that the level of logging is dependent on your license level.

- **Logging Capabilities** – Logs from M365 can also be ingested into a SIEM to retain data for longer periods of time.

Conditional Access Policies – These policies are able to limit access to M365 by placing certain restrictions on accounts such as, geolocation, preventing legacy protocols.

Data Loss Prevention (DLP) Software – DLP can be used by organisations to track data throughout the organisation. This can be used by organisations to track which files are being opened and by who. The process of implementing DLP can be long and expensive, however, it will help to quickly determine what documents have been accessed by any malicious entities.

Alert Policies – These are designed to categorise alerts that are triggered by a policy within M365. These rules can range across all facets of M365 and can be used to indicate whether inbox rule creation, additional MFA devices,

- **MFA Modification** – As we have shown earlier, Threat Actors are able to modify which devices are registered to the account for MFA.
- **Inbox Rules** – If a Threat Actor is able to steal a token, it is likely that the Threat Actor will create an Inbox rule to sift emails into other folders. This requires, at a minimum an E1/F1/G1 license.
- **Device Enrolment** – Depending on the permissions the Threat Actor has access to when they have access, the Threat Actor may enrol their own device to Azure Active Directory to bypass conditional access rules.
- **Microsoft Entra Security Reports** – There are a number of reports available to identify Risky or Anomaly sign ins to an account, some of these reports are only available to Microsoft Entra ID P2 customer.

It is important to note that no single option outlined above is a silver bullet that will render your environment invulnerable to token theft activities, but a combination of such options contributes to a layered defence in depth model that will either prevent it from happening or alert you to it happening when it does.

Phishing Resistant Authenticators

The use of Phishing resistance authentication methods can help prevent a Threat Actor from gaining access to an account. These work by using either something you have or something you are. For the something you have, this is often a smart card or a USB dongle. For the something you are this is usually some biometric, a face scan, fingerprint, iris scan, etc.

While at first glance something like the Microsoft Authenticator (when setup correctly) or the relatively new passkey systems that Microsoft, Google, and other providers are rolling out might appear to be resistant to token theft, and to a degree they are, the MFA part is resistant, but the token is still able to be compromised by the Threat Actor. This is because the challenge to verify the identity of the user is happening on a separate device, in a separate session to the one the attacker is trying to steal the token for.

To prevent token theft the challenge needs to happen in the same session as the login. This is where the use of Fast Identity Online (FIDO) 2.0 devices comes into play, these work by doing the challenge in the same session that you're logging in as. They work by using a combination of public and private key encryption that is setup with the FIDO device is enrolled, as well as information stored on the FIDO device about what that key is for i.e. the URL for the website you've just enrolled it for.

The public and private key encryption works to ensure that subsequent authentication isn't intercepted and compromised by a Threat Actor but doesn't stop them from sitting in the middle of the session because they can just pass the authentication through without altering it. This is where the information stored on the FIDO at the time FIDO device is enrolled is used. When trying to authenticate the FIDO device will see that the URL that you're trying to sign in to isn't the URL that the FIDO was setup with and not allow the authentication to happen preventing the token from being generated and then stolen.

Device Certificates and Smart Cards are also other examples of phishing resistant authentication methods.

Since accounts with phishing resistant authentication is much harder to compromise that those without. Where possible it should be implemented on Administrator and other high privileged accounts.

To prevent token theft the challenge needs to happen in the same session as the login.

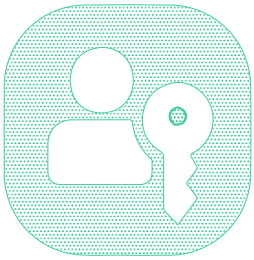

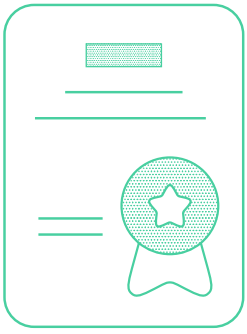

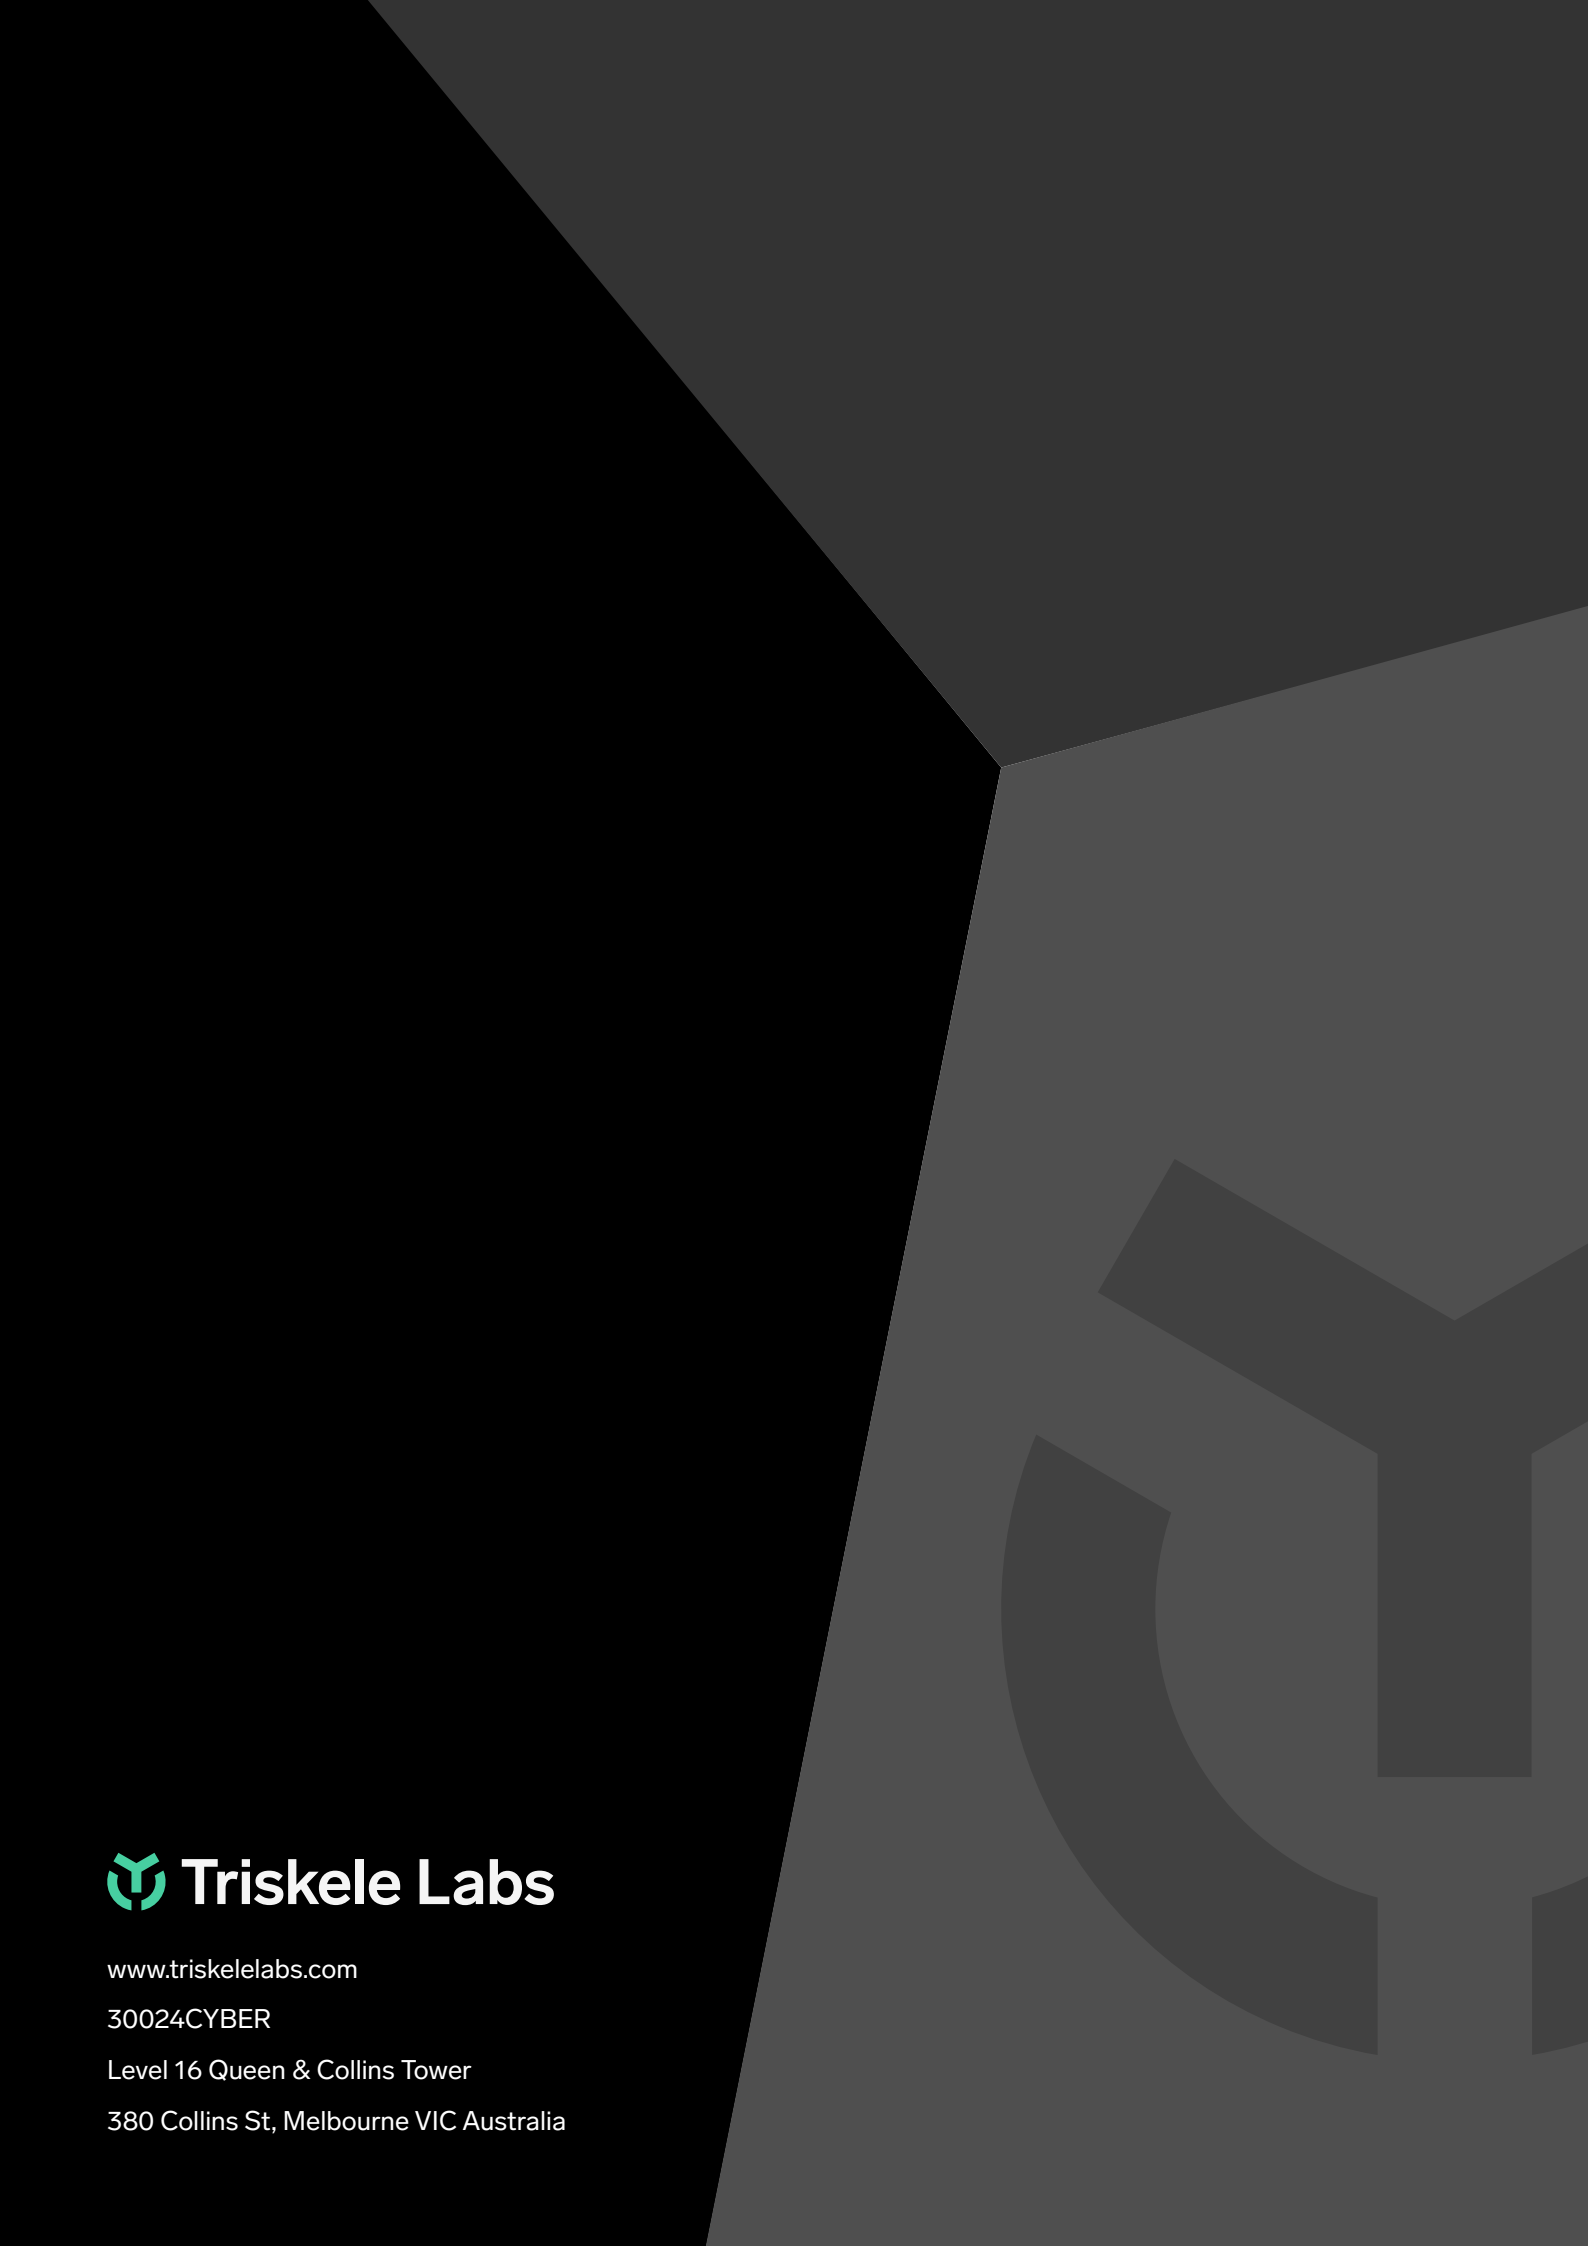

Triskele Labs

www.triskelelabs.com

30024CYBER

Level 16 Queen & Collins Tower

380 Collins St, Melbourne VIC Australia