

Phishing in the deep blue Azure

Jacob Torrey
@jacob@mountaincommunity.co



Agenda

- Introduction & background
- Addressing the visibility gap
- The token
- Data and how it's changed our minds
 - Common threat data
 - Interesting data
 - Useful data
 - Data worth sharing
- Sharing data
- Caveats
- Conclusions
- Q&A

The team



Jacob
Head of Labs



Casey
Sr. Security
Engineer



Jay
SW Engineer



Max
FE Engineer



Gareth
CS Engineer

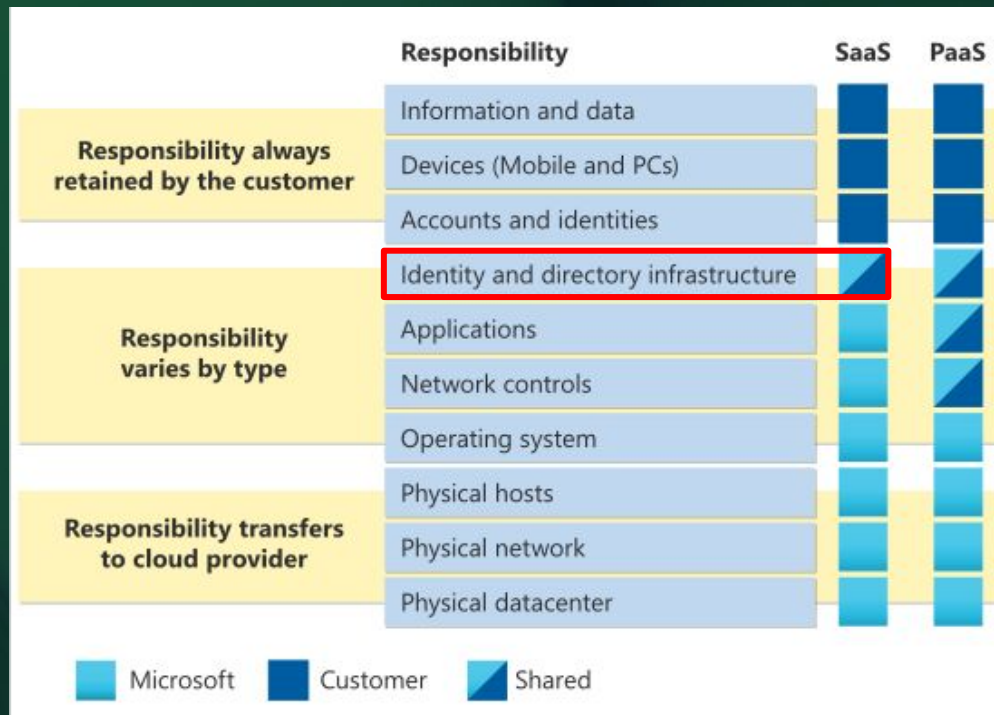
Thinkst Labs

- Thinkst helps customers know... when it matters
- Labs is the research group within Thinkst
- Labs also publishes a quarterly research review:
ThinkstScapes
 - We spend a lot of time on making this a good read...
 - And we give it away for free: thinkst.com/ts



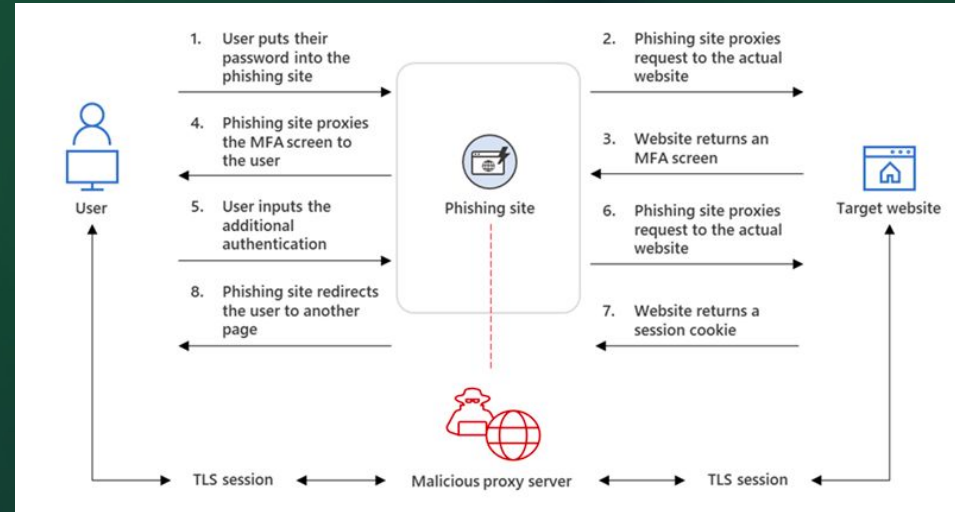
The visibility gap

- Deception engineering is a nascent field, and it's *hard*
- The SaaS-ification of business critical infrastructure means there are more blind spots at the seams of the shared responsibility model
 - Attackers exploit these seams where context is limited



Background: Adversary-in-the-middle phishing

- AitM is an evolution of cloning a victim website
- The attacker acts as a reverse proxy between the victim site and the victim user
- AitM can defeat most MFA
- Massive growth in AitM phishing
 - > 50% YoY growth 2021-2022
 - Kroll (IR firm) reports that in 90% of their investigations Q2&3'23 MFA was in place but M365 sessions were still stolen



Target for this project: Entra ID

- Where could we build a lightweight sensor that provides alerts on detected badness?
- Azure Entra ID (or more broadly `login.microsoftonline.com`) is the landing page for many organizations
 - Microsoft doesn't have all the context of who's expected to be logging in, and how for each tenant
 - The tenant owner doesn't see the raw telemetry pre-auth
- The Entra ID Canarytoken is one small step towards improved visibility

The token

- Entra ID allows for tenants to customize their login page
 - Most importantly CSS
- The body element is completely covered by other content, so we can safely change its background to a `url()` that points to our serverless functionality
- The Entra ID login page specifies a referrer-policy
- Serverless function checks to see if the `Referer` matches a Microsoft domain



jacobdev@scs[REDACTED].onmicrosoft.com

Permissions requested

Review for your organization



This app would like to:

- ✓ Read and write all applications
- ✓ Read and write organization information
- ✓ Sign in and read user profile

If you accept, this app will get access to the specified resources for all users in your organization. No one else will be prompted to review these permissions.


Accepting these permissions means that you allow this app to use your data as specified in their [terms of service](#) and [privacy statement](#). You can change these permissions at <https://myapps.microsoft.com>. [Show details](#)

Does this app look suspicious? [Report it here](#)

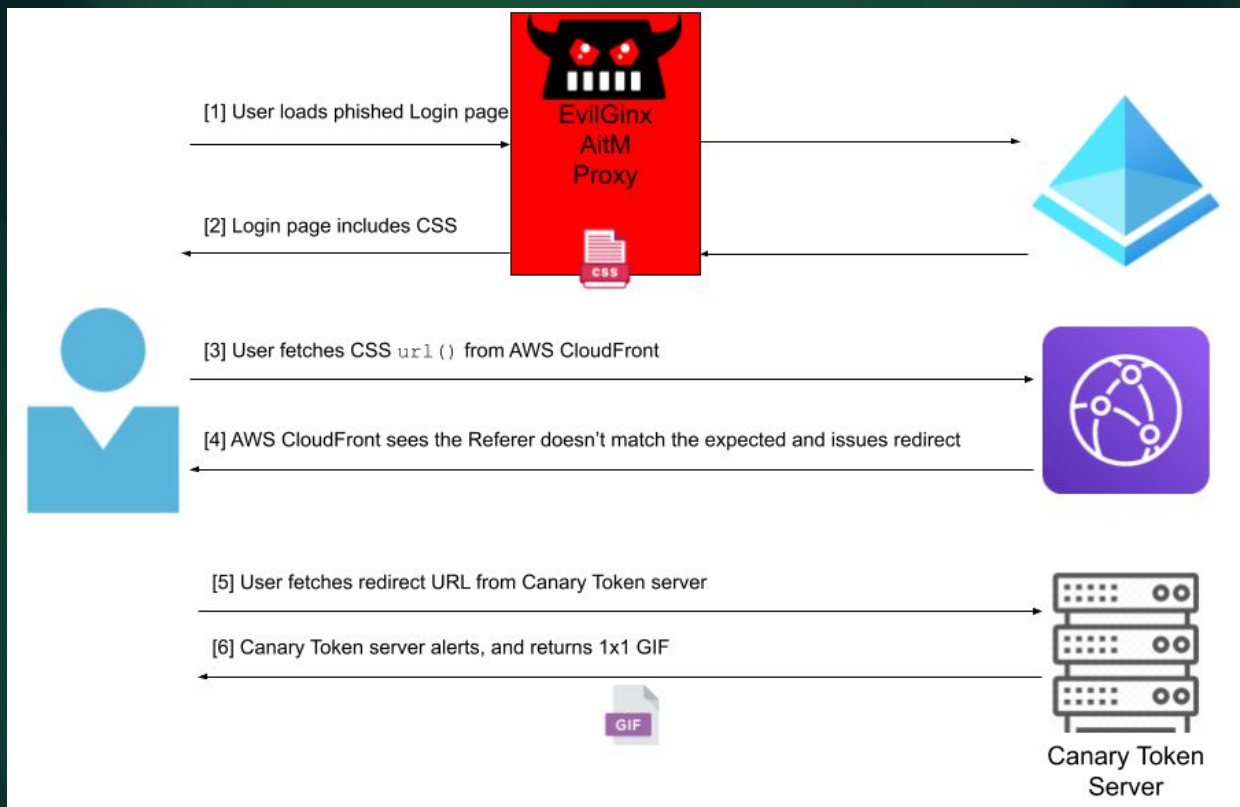
Cancel


Accept

Happy case



Unhappy case





Canarytoken triggered

ALERT

A CSS cloned website Canarytoken has been triggered by the Source IP 3. [REDACTED]

Basic Details:

Channel	HTTP
Time	2024-01-24 21:57:03.108330
Canarytoken	h6jf84z88jbpfc1at5n5blc1
Token reminder	Blog and personal site
Token type	CSS cloned website
Source IP	3. [REDACTED]
User-agent	Mozilla/5.0 (Windows NT 10.0; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/120.0.0.0 Safari/537.36
Referer	https://test.tlsdebug.com/

Canarytoken Management Details:

Manage this Canarytoken [here](#)

More info on this token [here](#)

Powered by [Thinkst Canary](#)

Did you know some of the best
security teams in the world run
Thinkst Canary?

Find out why



Caught you!

Now...
Show me the data

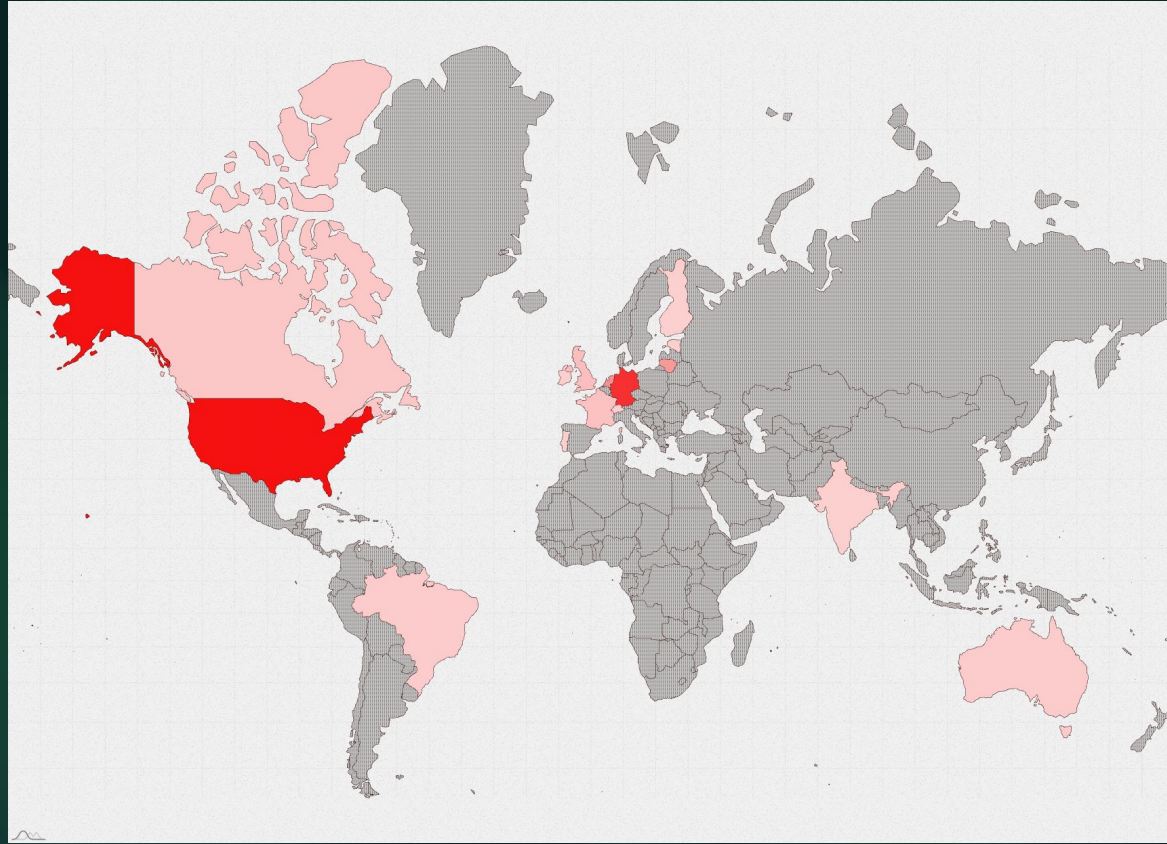


Our take on data sharing

- We're not a threat intelligence firm
- Threat sharing data is often: interesting **but not as often:** useful
 - The data was useful to the tenant owner, but others?
- Not worth sharing data that's only interesting
 - Until recently we didn't think any of our data was useful
 - **E.g., 85%** of the domains we alert on are seen only for a single day


Too often you see these
types of charts

Interesting? Infrastructure analysis



Interesting? Infrastructure analysis



Phishing certificates issued





Then you get the
feel-good data

From 0 to millions for some pizzas

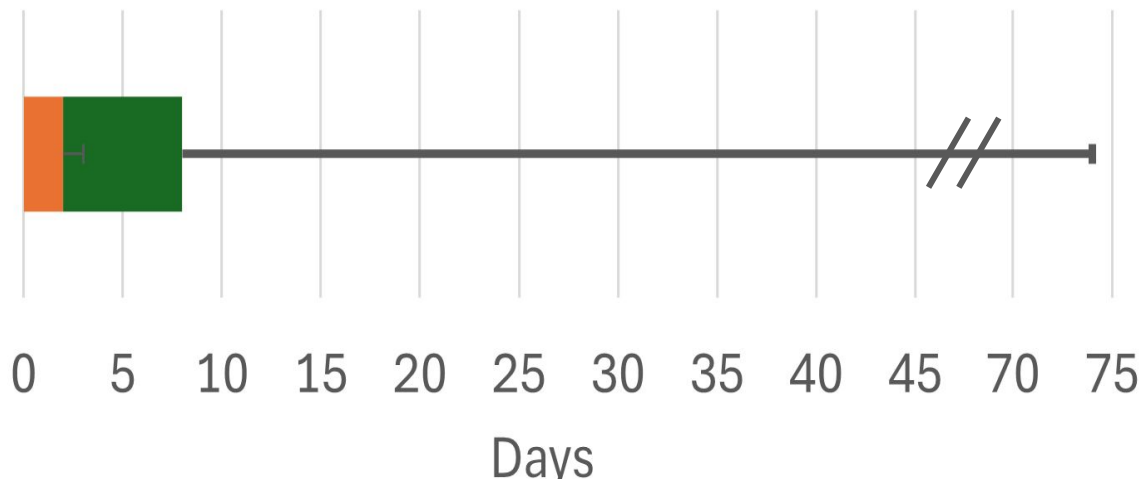


*Data as of 10/10, only including the open-source canarytokens.org


Alert analysis

~45% detected in first day after certificate issuance!
~75% in first week!

Time from certificate creation to first alert



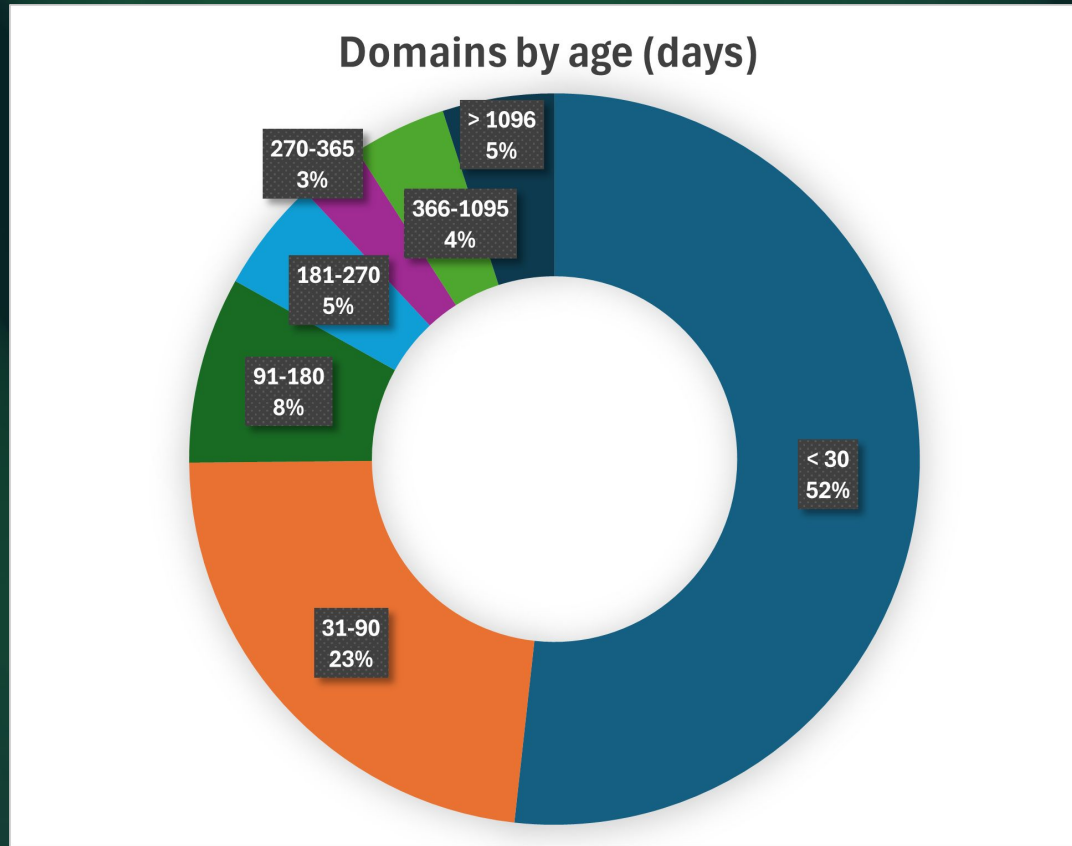
Time from domain registration to first alert



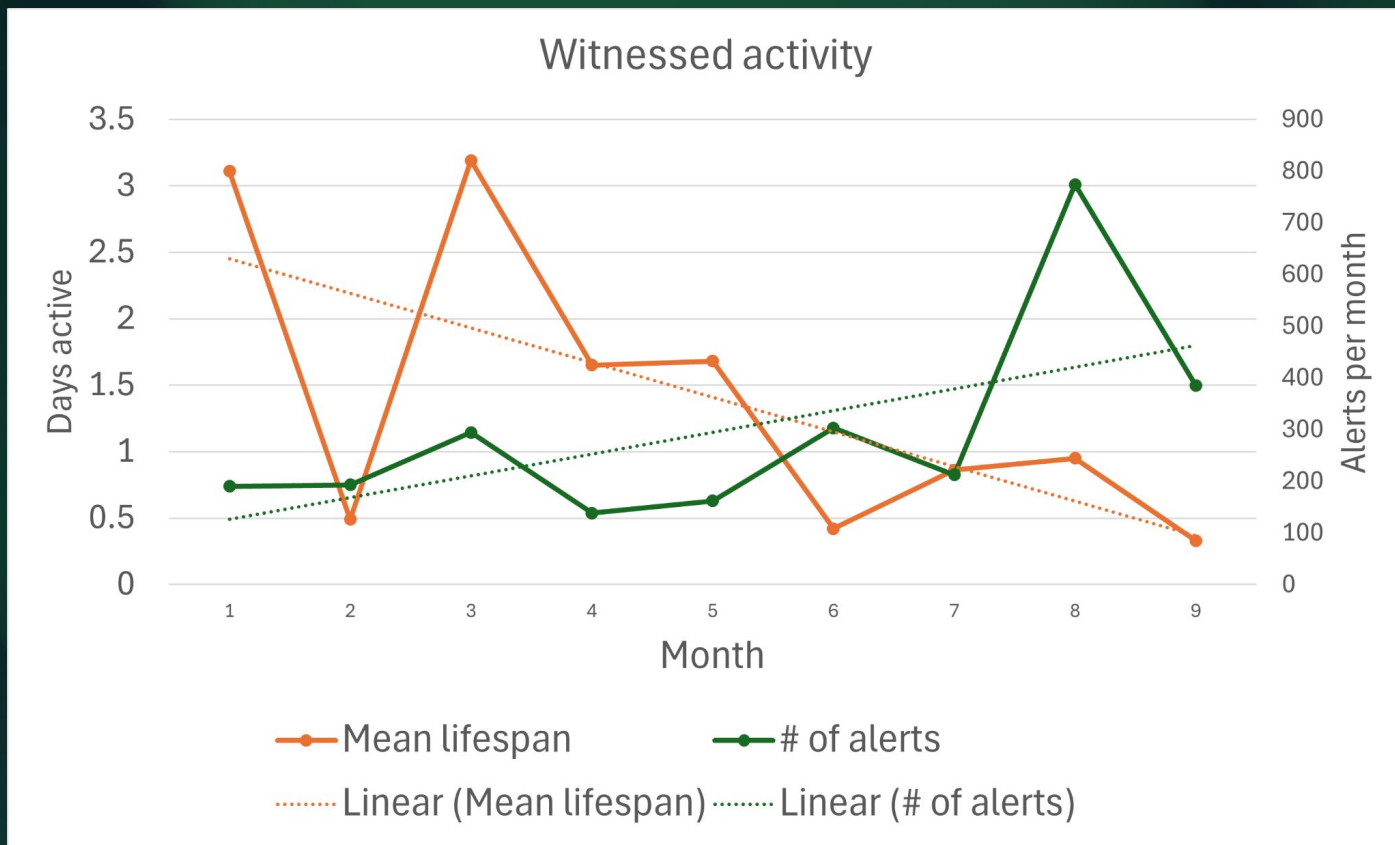
15% detected in first day of domain registration!

Then the interesting data

Split of new and "seasoned" domains



Our view into malicious activity



Early on: An interesting case

- “Splash damage”: One domain -> multiple tenants
- Example:
 - Domain was <biotech firm name> + bio.com
 - <biotech firm name>.com is the real domain
 - All the hallmarks of AitM phishing
 - Apex domain redirects to example.com
 - Recently registered domain and Let’s Encrypt certificate
 - And yet we see this domain phishing another tenant
- Interesting since the phishing tools are incredibly touchy to prevent getting blocked



Now we see
multi-tenant alerts
all the time...

What does this mean?

- A change in attack



An indication that our data is useful



Leads us to...
useful data?

Our thinking has changed

- We have a wider aperture
 - More cross-tenant hit visibility
- See more badness
 - While only a minority of domains last > 1 da
- Seen data for longer
 - Know the dangerous corners of the internet
 - Top 3: .workers.dev, **.azurewebsites.net**, .web.core.windows.net



Why care about our data?

- We see only successful phishing – irrespective of methods
- We see data from behind the defenses
- We see data early
 - The number of alerts we see within a day of creation means we must be catching attackers testing...

Sharing the data!

How/where to share?

- Directly with those hosting these sites?
 - A lot of work, but would end up with the best resolution









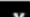


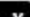
















Hello,

This is a follow up regarding the abusive content or activity report that you submitted to AWS. After our investigation, we are unable to verify your claim. Please provide additional details to assist us in further investigation.



Now starting to flow some domains into MISP

Azure Entra ID phishing

Event ID	264046
UUID	d383aff1-d9c0-49e4-8fe8-45efff3b2b53  
Creator org	thinkst.com
Creator user	Jacob@thinkst.com
Protected Event (experimental) 	 Event is in unprotected mode.
Tags	 tlp:clear    PAP:CLEAR    Phishing    circl:incident-classification="phishing"    phishing:techniques="fake-website    Credential Phishing      

- Domain and full ORE
- IPs
- LookyLoo analysis

Caveats

Adversaries adapt

- Already some red teamers recognize the risks of the Referrer header
 - Inject a referrer-policy header from AitM to provide none
 - Browser-in-the-middle that runs e.g., a Chromium browser and sends only the pixel data
- We currently **do not** alert on blank referrers
 - The tooling as available doesn't change the referrer policy
- We *could* change that behavior easily
 - We constantly are evaluating the false positive risk, and how noisy it would be
 - Consistently we see about 0.05% (½ of 1/10th of a percent) of hits have a blank referrer
- Also... Sometimes the good guys trigger alerts
 - Microsoft SmartScreen, and other phishing block-lists will revisit/retrigger to see if adversary infrastructure is still active

Band-aids shouldn't stop aiming to a gold standard

- Authentication providers should be pushing **phishing-resistant** MFA
- Conditional Access shouldn't require disabling the industry-standard security settings
 - Tenant owners should be able to have both

Conclusions

- AitM tooling is making it easier to steal sessions and credentials, even with MFA
 - Entra ID is a popular target for AitM phishing
- SaaS outsourcing reduces visibility, making detection engineering harder
- The Entra ID token offers a sensor to help organizations detect AitM against their tenants before the victim even logs in
 - Free @ canarytokens.org!
 - Gaining popularity, more popularity = more data on attackers and faster response
 - Data starting to flow to MISP
- Still need to move people away from phishable credentials
- Consider visibility costs with outsourcing

Thank you!

Q&A



ThinkstScapes



CanaryTokens