Sweet Deception:

Deploying Honey Tokens in Microsoft 365

Ryan O’Donnell
in/odonnell-ryan/

Agenda

1. Background
2. Design
3. Results
4. Conclusion

Disclaimer

The opinions and content expressed are my own and do not reflect my employer’s views. The content is based on my personal research and experience. Mention of specific tools, technologies, or companies does not imply endorsement by my employer. I do not assume any liability or responsibility for any errors or omissions in the content or for any losses, damages, or injuries resulting from the use of the information provided. The information, opinions, and materials provided in this presentation are for educational purposes only and should not be considered as professional advice. My employer has no involvement in the creation, review, or endorsement of the content in this presentation.

whoami

- Ryan O’Donnell
- Senior Security Engineer @Microsoft
- Co-Founder @Vortacity
- Penetration Testing, Purple Teaming, Red Teaming, DFIR
- OSCP, OSEP, GCFA, GREM

Cybersecurity Reality

Florian Roth ⚡

@cyb3rops

We're seeing a clear trend: attackers are bypassing the endpoint entirely. Not just avoiding traditional EDR-monitored systems by pivoting to embedded and edge devices, but now also operating purely in the cloud. No shell, no malware, no persistence on the endpoint. Just an OAuth token and full access to whatever's in the victim's Microsoft 365, Google Workspace, or AWS console.

Florian Roth: @cyb3rops

Current Trends

- Cloud intrusions **increased 136%**
- China based cloud intrusions **up 40%**
- 81% of intrusions were **malware-free**

Crowdstrike 2025 Threat Hunting Report

Current Trends

- Cloud intrusions **increased 136%**
- China based cloud intrusions **up 40%**
- 81% of intrusions were **malware-free**

Crowdstrike 2025 Threat Hunting Report

Proposal

Canary Tokens can provide effective **early alerting** on attacker **post-compromise** activity.

Background

Deception

- **Honeypots**
  - Decoy systems
  - Log behavior
- **Canary Tokens**
  - Discrete data elements
  - Function as tripwires

How Are Breaches Discovered

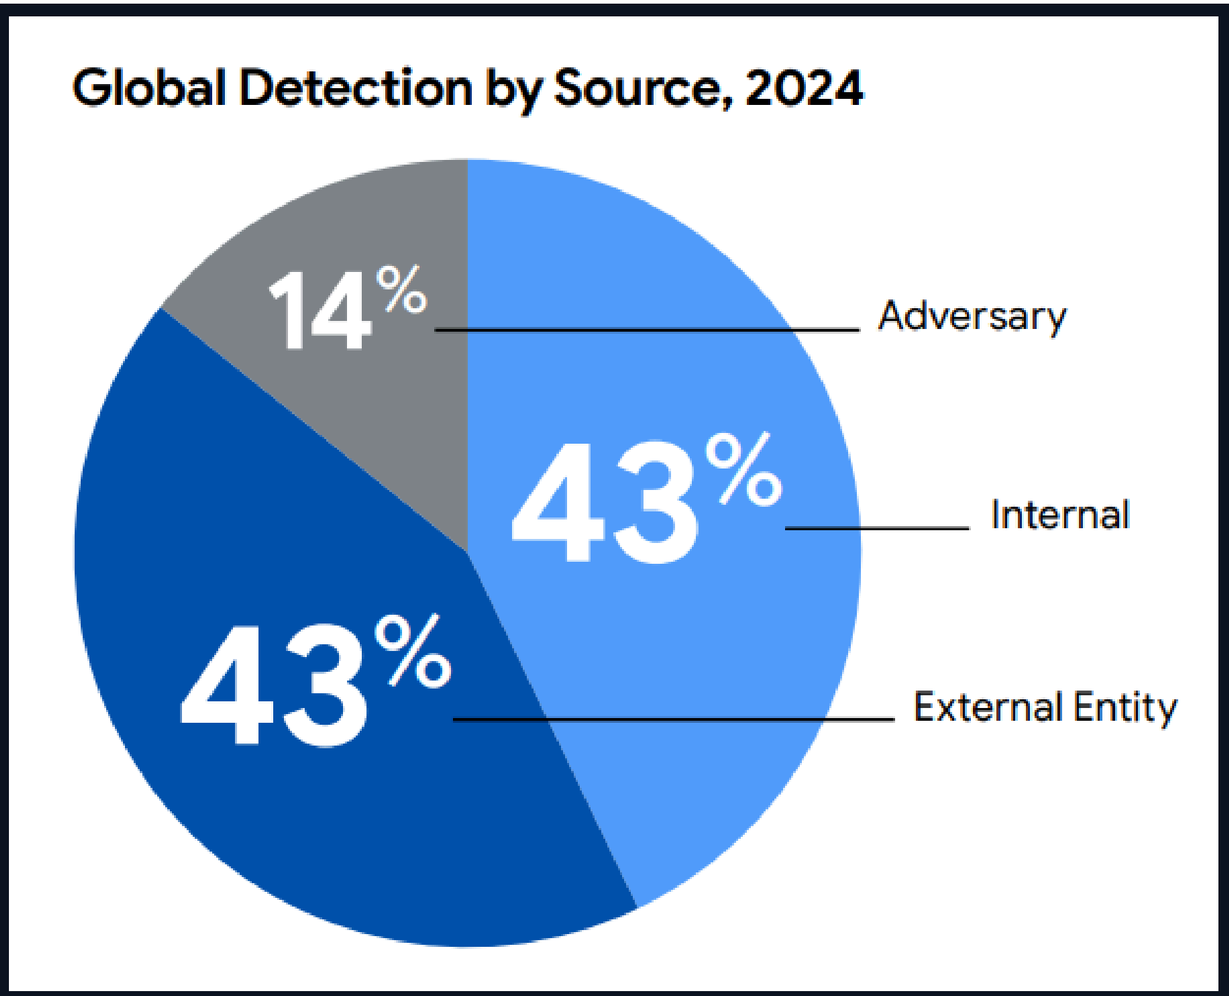

Mandiant M-Trends 2025 Report

Dwell Time

Median Dwell Time by Detection Source, 2024

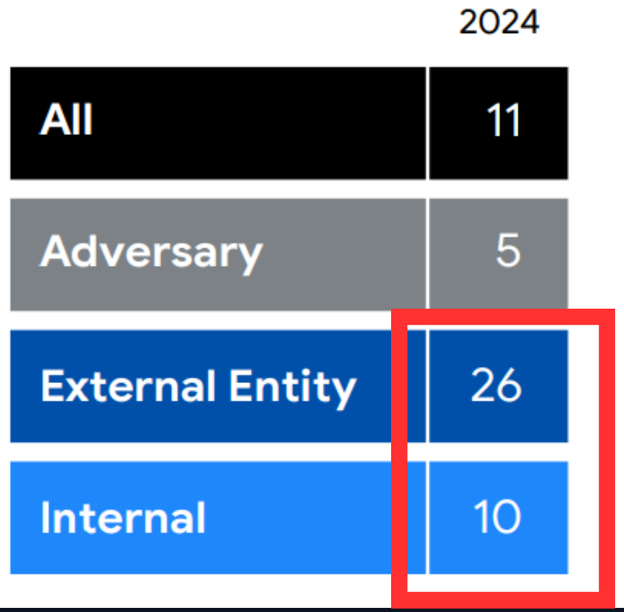

The median adversary notification time was just five days, while external partners notified in a median of 26 days. This discrepancy is not surprising given that the vast majority of adversary notifications originate from extortion actors who benefit from monetizing intrusions quickly.

Mandiant M-Trends 2025 Report

Deployment Strategies

Criteria

Attractive

High Fidelity

Non-Invasive

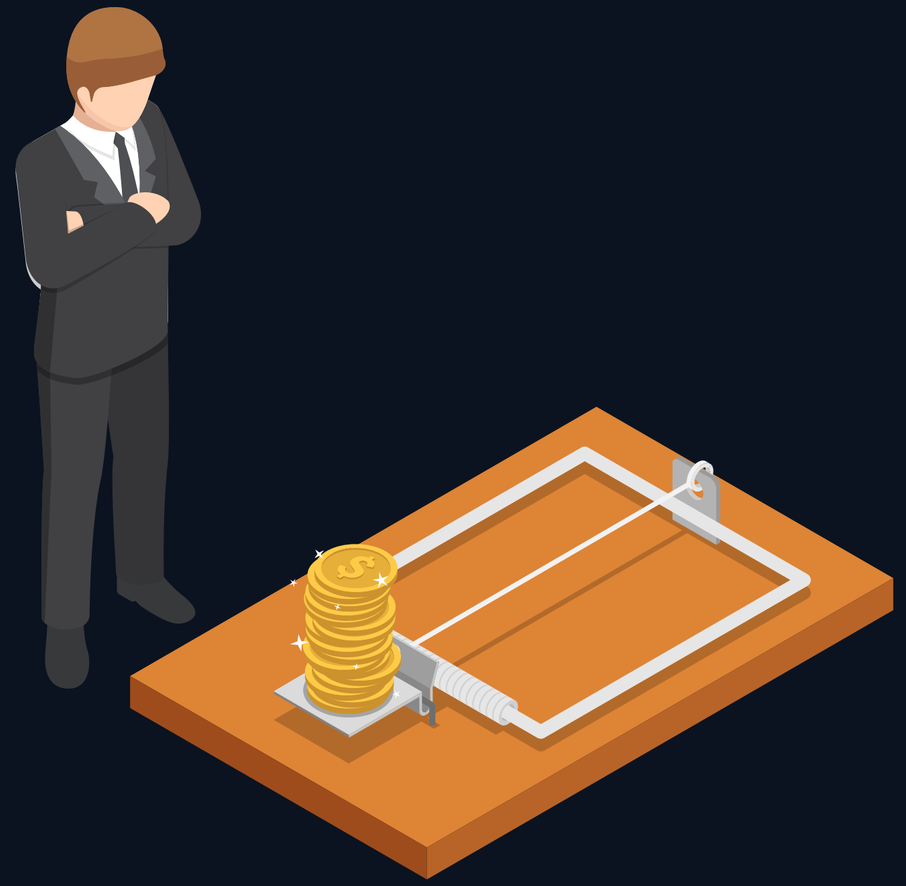

Post-Compromise

After gaining initial access, Void Blizzard abuses legitimate cloud APIs, such as Exchange Online and Microsoft Graph, to enumerate users' mailboxes, including any shared mailboxes, and cloud-hosted files. Once accounts are successfully compromised, the actor likely automates the bulk collection of cloud-hosted data (primarily email and files) and any mailboxes or file shares that the compromised user can access, which can include mailboxes and folders belonging to other users who have granted other users read permissions.

MSTIC: Storm-2372, Cookie Theft To BEC, Void Blizzard

Crown Jewels

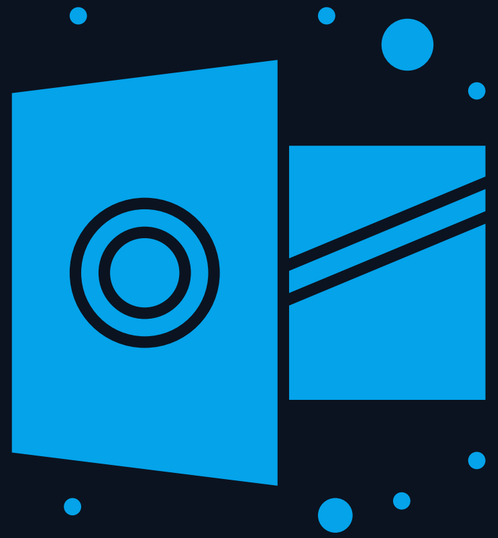

Real-World Implementation

Existing Approach

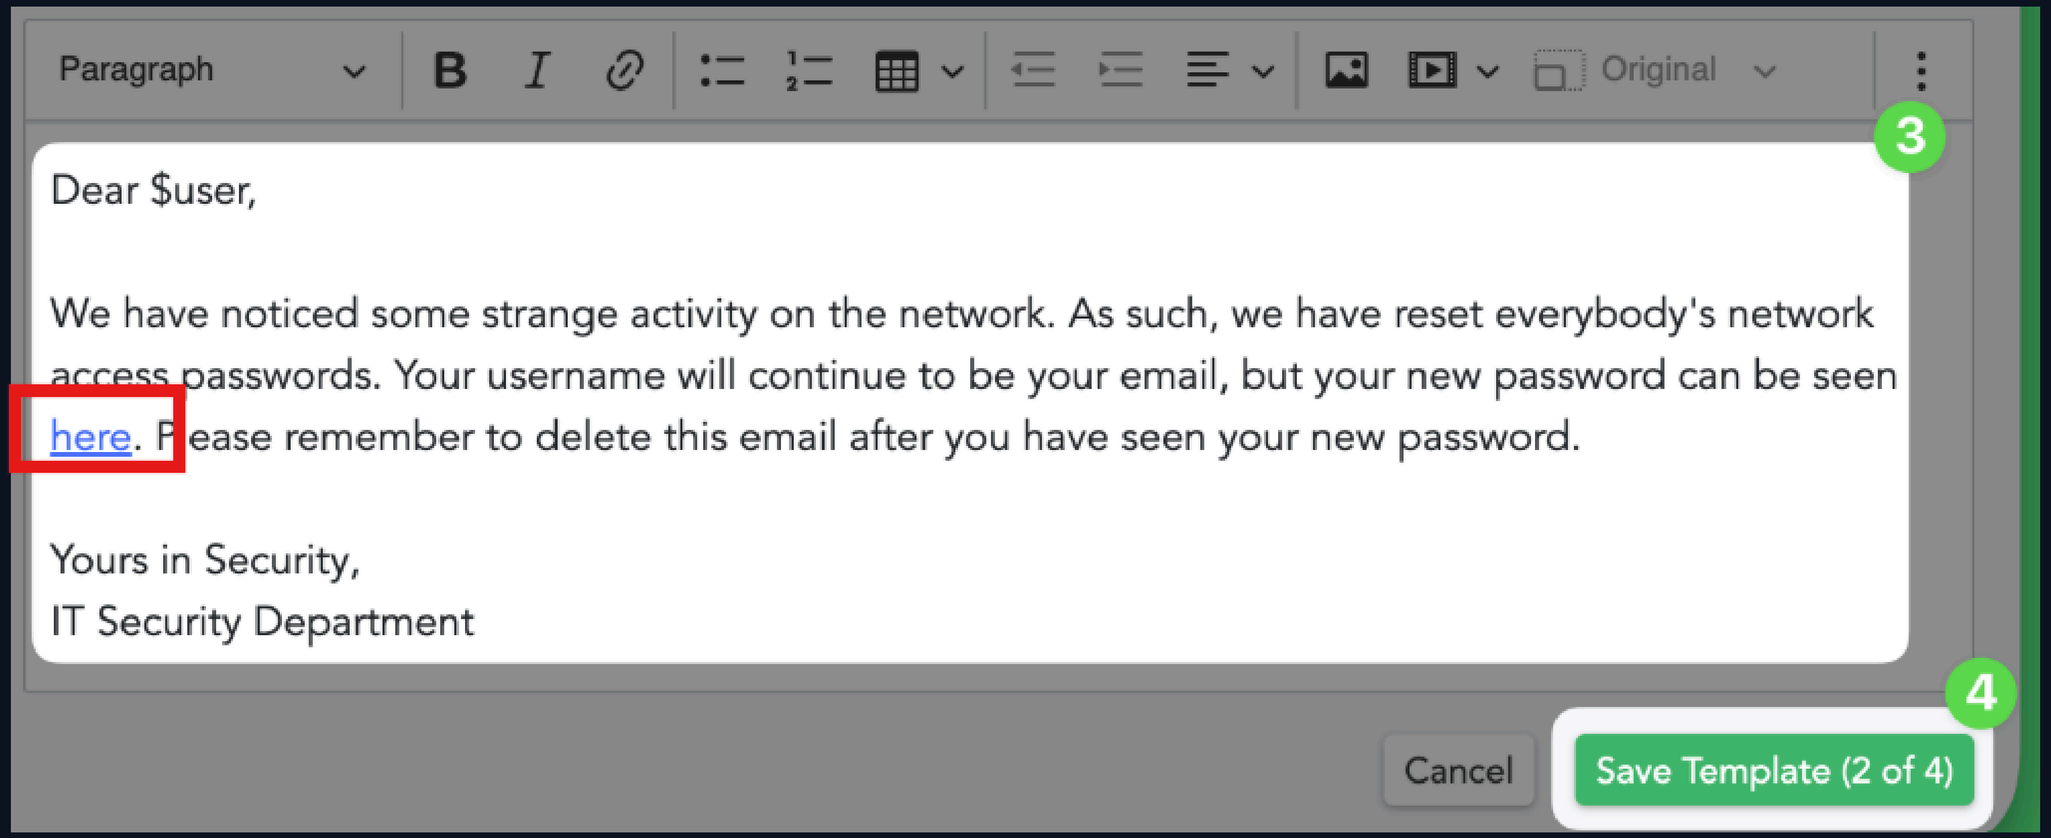

Lure

- Creds: username, password, admin
- Finance: invoice, wire transfer, SWIFT
- IT: API keys, root, tokens
- RMM: anydesk, screenconnect
- HR: payroll, W2, workday

Research

Hidden mail folders

The default value of the `isHidden` property is `false`. You can set **isHidden** only once when creating the mailFolder. You can't update the property using a PATCH operation. To change the **isHidden** property of a folder, delete the existing folder and create a new one with the desired value.

Hidden mail folders support all operations that are supported by a regular mail folder.

Soteria: Hidden Mailbox Folders

Search

```
PS /opt/tools/GraphRunner> Invoke-SearchMailbox -Tokens $tokens -SearchTerm "credentials"
[*] Using the provided access tokens.
[*] Found 9 matches for search term credentials
Subject: Info | Sender: finance@azpurple.com | Receivers: Adele Vance | Date: 06/28/2025 17:59:30
| Message Preview: Credentials Password Information ...
================================================================================
Subject: Dev Environment | Sender: /O=EXCHANGELABS/OU=EXCHANGE ADMINISTRATIVE GROUP (FYDIBOHF23SPD
LT)/CN=RECIPIENTS/CN=E286EAE0D96F4FC89B9FD61675ED286F-903988D0-C8 | Receivers: ippsec | Date: 03/0
9/2025 19:31:32 | Message Preview: ... Below are the temporary credentials to the developer enviro
nment. User: mark Pass: MDRisTh3best! Host: 192.168.1.12  Thanks, Adele ...
================================================================================
[*] Do you want to download these emails and their attachments? (Yes/No)
Yes
```

Monitor

- **MailItemsAccessed** operation, UAL


```
"ClientRequestId":  "7c0594c8-397c-4f8d-88cc-ffd0c3b366e8",
"Id":  "RgAAAAC5oCuB0sfIRpW3nYYIkzQeBwDVYET+o7QuT6cHeqkYpuL7AAAAAAEM
"ImmutableId":  "LgAAAAAdhAMRqmYRzZvIAKoAL8RaDQDVYET+o7QuT6cHeqkYpuL
"InternetMessageId":  "\u003cCY8PR10MB665782D7B132D0C5692E0DEFC9D22@
"Sensitivity":  "defa4170-0d19-0005-0004-bc88714345d2",
"SizeInBytes":  38085
```

```
"AADSessionId":  "00401e29-2d0f-1197-87ce-6b4235eb9041",
"APIId":  "00000003-0000-0000-c000-000000000000",
"ClientAppId":  "d3590ed6-52b3-4102-aeff-aad2292ab01c",
"IssuedAtTime":  "2025-04-20T15:59:18",
"UniqueTokenId":  "0jpPuyxmikOdByzOaIcXAA"
```

Process

1. Create **hidden folder**
2. Send an email with **keywords**
3. Move email to hidden folder
4. Monitor **MIA** events
5. Alert when email is **accessed**

Test in Production

- 50 employees
- Engineers/sales
- Cloud-native
- Outlook Canaries
- **Subset** of users

Results


10
8
6
4
2
0
Number of Events
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28
Time (Days)

Results

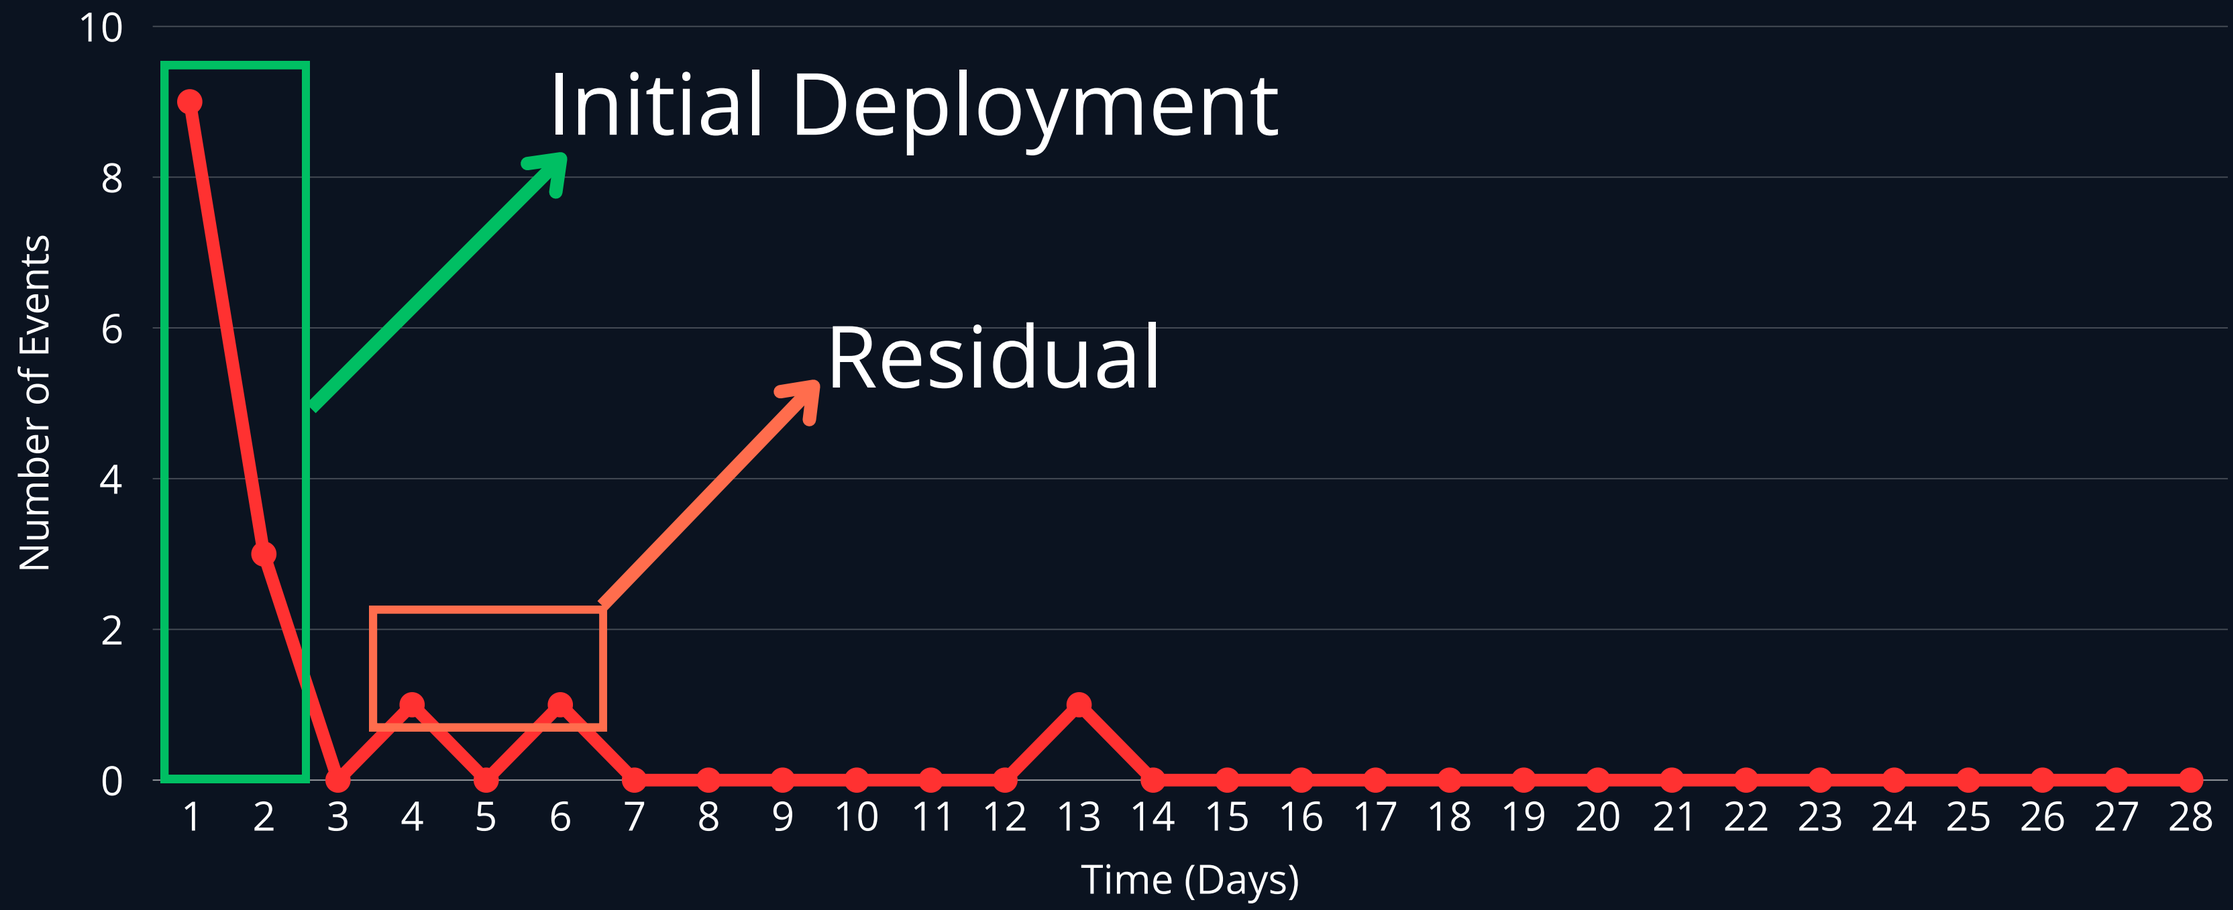
Initial Deployment
Residual
Number of Events
Time (Days)
0
2
4
6
8
10
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28

Suspect

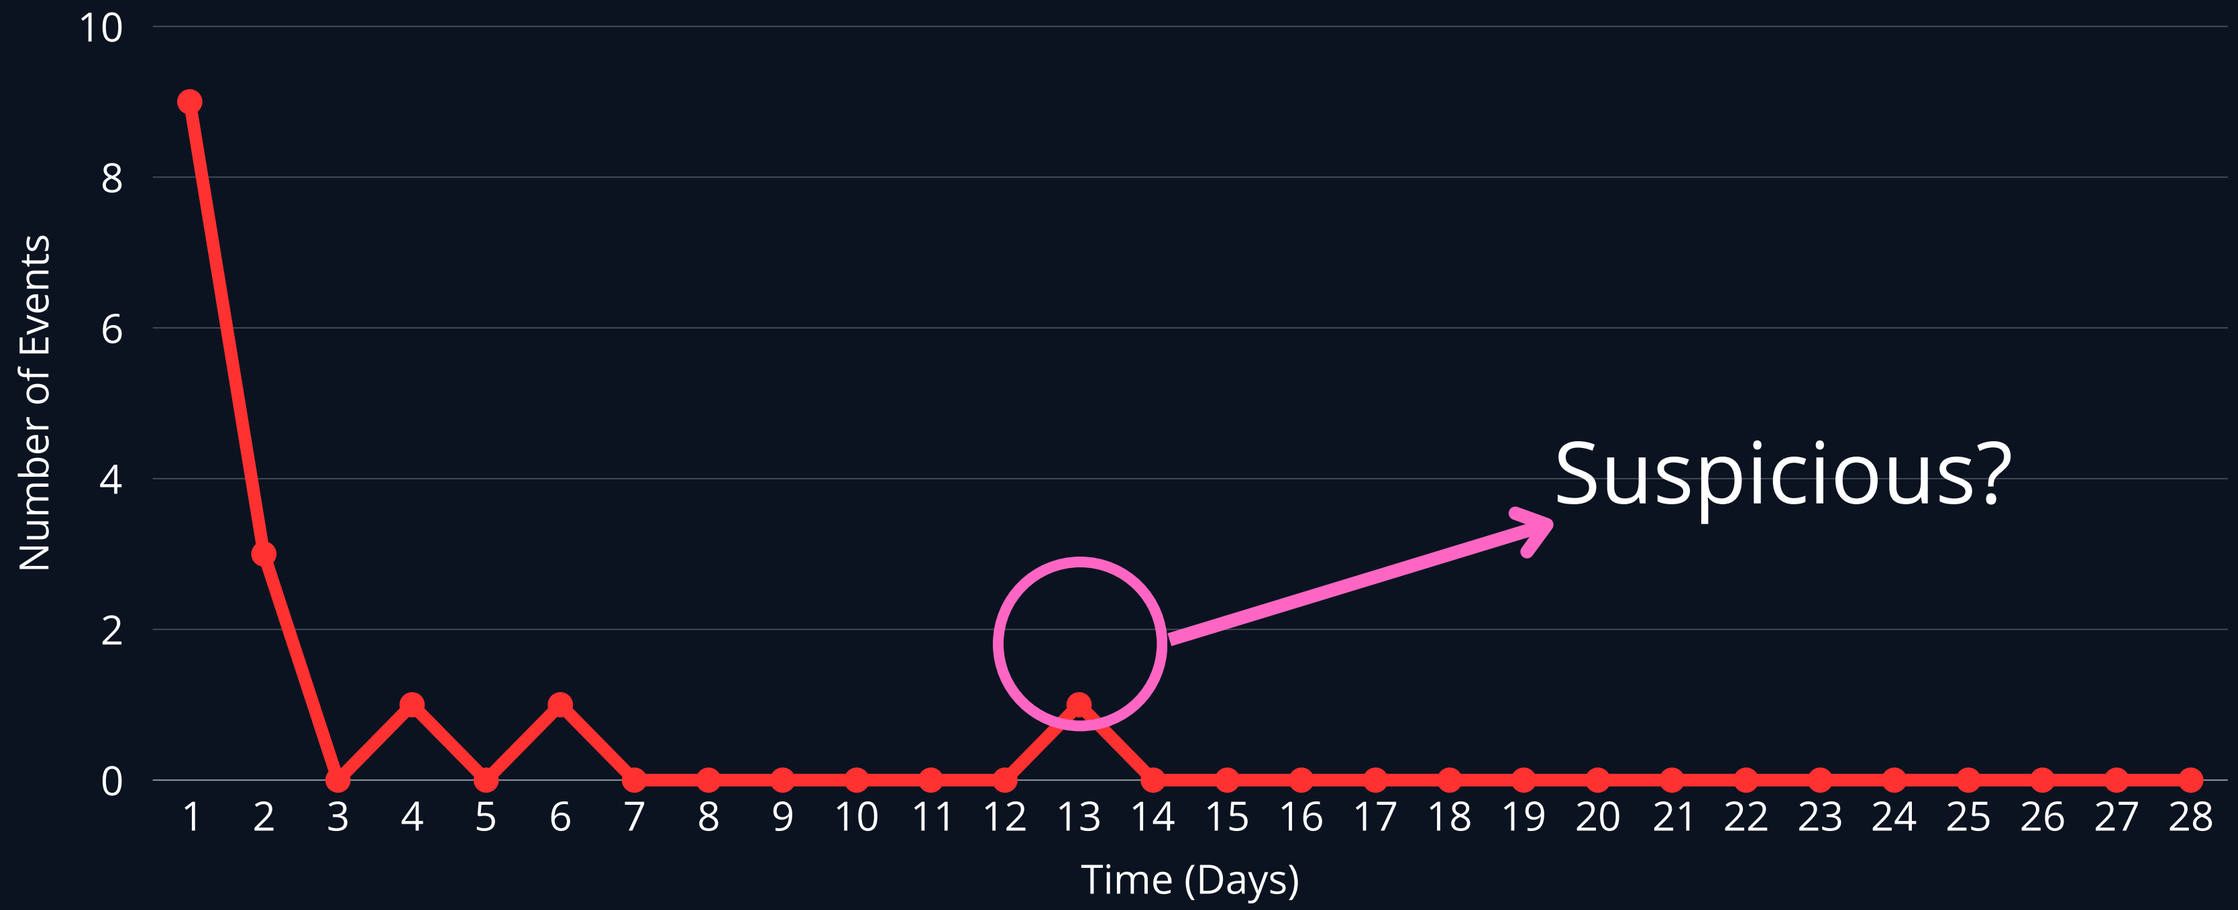
10
8
6
4
2
0
Number of Events
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28
Time (Days)
Suspicious?

Suspect

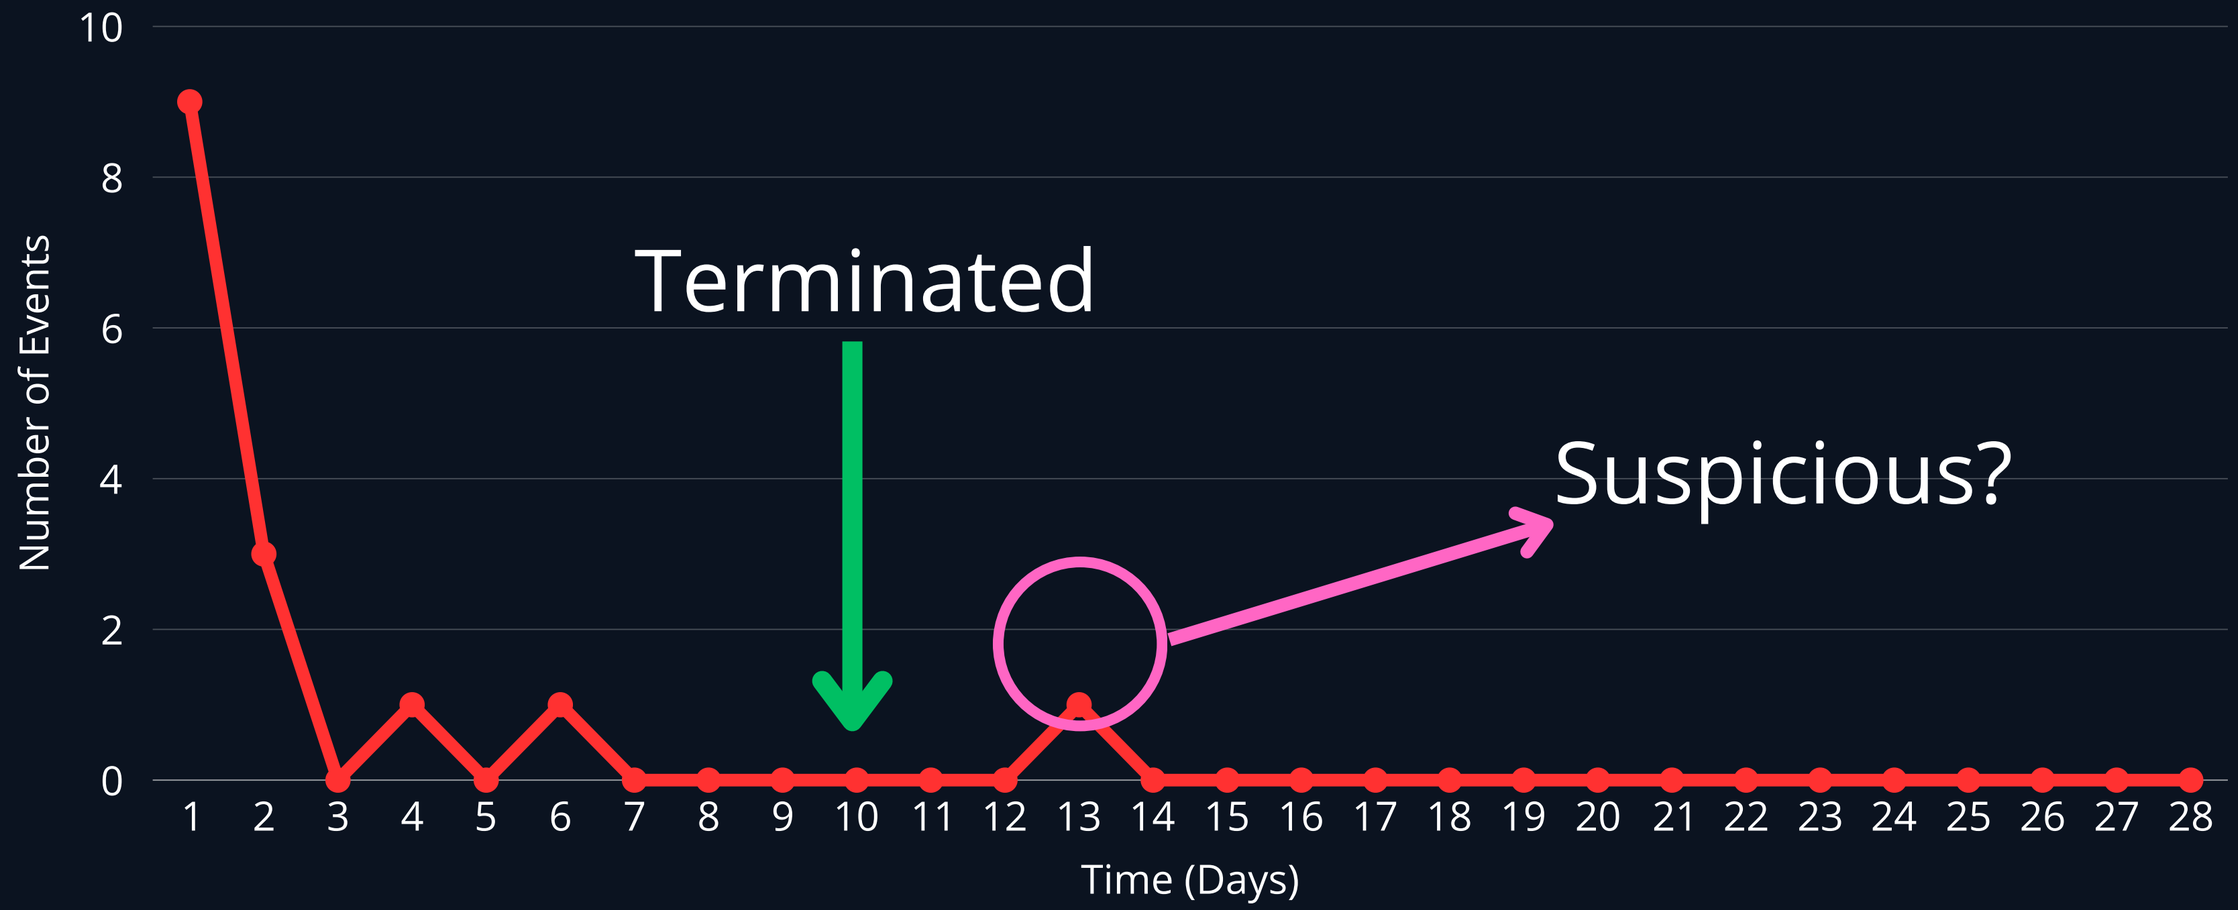
10
8
6
4
2
0
Number of Events
Terminated
Suspicious?
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28
Time (Days)

Conclusion

Evolution

And that’s the shift: attackers aren’t hacking computers anymore. They’re hacking trust relationships, identities, and APIs. The whole idea of detection and response needs to evolve with that. Otherwise, we’re securing the hell out of endpoints while attackers happily fish through mailboxes and cloud shares from halfway across the planet.

Florian Roth: @cyb3rops

Questions?

Contact: in/odonnell-ryan/

Twitter/X: **@odiesec**