New Attack ‘Clones’ and Abuses Your Unique Online ID via Browser Fingerprinting

Originally published at https://www.unite.ai/new-attack-clones-and-abuses-your-unique-online-id-via-browser-fingerprinting/

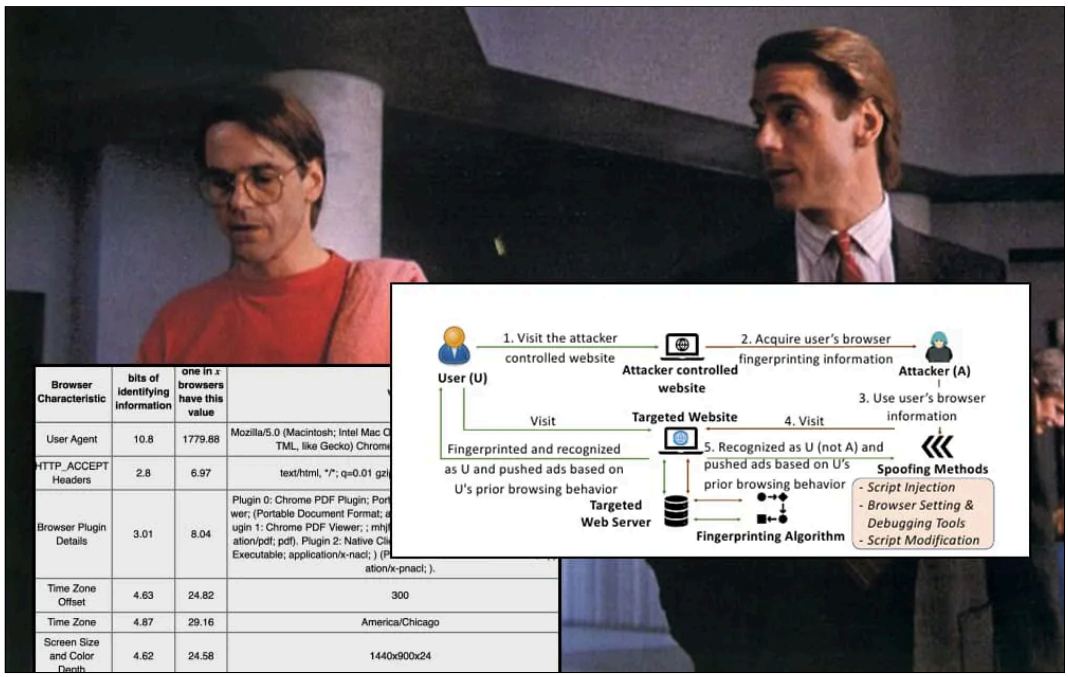

Researchers have developed a method to copy the characteristics of a victim’s web browser using browser fingerprinting techniques, and thereafter ‘impersonate’ the victim.

The technique has multiple security implications: the attacker can carry out damaging or even illegal online activities, with the ‘record’ of those activities attributed to the user; and two-factor authentication defenses can be compromised, because an authenticating site believes that the user has been successfully recognized, based on the stolen browser fingerprint profile

Additionally, the attacker’s ‘shadow clone’ can visit sites that change the kind of ads delivered to that user profile, meaning that the user will start receiving advertising content unrelated to their actual browsing activities. Furthermore, the attacker can infer much about the victim based on the way other (oblivious) websites respond to the spoofed browser ID.

The paper is titled *Gummy Browsers: Targeted Browser Spoofing against State-of-the-Art Fingerprinting Techniques*, and comes from researchers at Texas A&M University and the University of Florida at Gainesville.

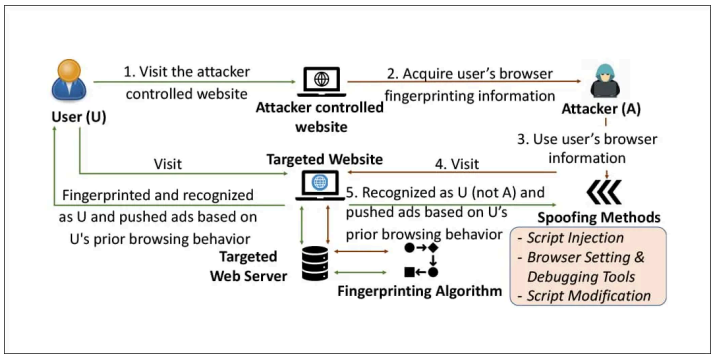

Overview of the Gummy Browsers methodology. Source: https://arxiv.org/pdf/2110.10129.pdf

Gummy Browsers

The eponymous ‘gummy browsers’ are cloned copies of the victim browser, named after the ‘Gummy Fingers’ attack reported in the early 2000s, which replicated victim's actual fingerprints with gelatin copies in order to bypass fingerprint ID systems.

The authors state:

> *‘The main goal of Gummy Browsers is to fool the web server into believing that a legitimate user is accessing its services so that it can learn sensitive information about the user (e.g., interests of the user based on the personalized ads), or circumvent various security schemes (e.g., authentication and fraud detection) that rely on the browser fingerprinting.’*

They continue:

> *'Unfortunately, we identify a significant threat vector against such linking algorithms. Specifically, we find that an attacker can capture and spoof the browser characteristics of a victim's browser, and hence can "present" its own browser as the victim's browser when connecting to a website.'*

The authors contend that the browser fingerprint cloning techniques they have developed threaten *'a devastating and lasting impact on the online privacy and security of the users'*.

In testing the system against two fingerprinting systems, FPStalker and the Electronic Frontier Foundation's Panopticlick, the authors found that their system was able to simulate the captured user information successfully nearly all the time, despite the system not accounting for several attributes, including TCP/IP stack fingerprinting, hardware sensors and DNS resolvers.

The authors also contend that the victim will be completely oblivious to the attack, making it difficult to circumvent.

Methodology

Browser fingerprinting profiles are generated by multiple factors of the way the user's web browser is configured. Ironically, many of the defenses designed to protect privacy, including installing adblocking extensions, can actually make a browser fingerprint more distinct and easier to target.

Browser fingerprinting does not depend on cookies or session data, but rather offers a largely unavoidable snapshot of the user's set-up to any domain that the user is browsing, if that domain is configured to exploit such information.

Away from overtly malicious practices, fingerprinting is typically used to target advertisements at users, for fraud detection, and for user authentication (one reason why adding extensions or making other core changes to your browser can cause sites to demand re-authentication, based on the fact that your browser profile has changed since your last visit).

The method proposed by the researchers only requires the victim to visit a website that's configured to record their browser fingerprint – a practice that a recent study estimated is prevalent on more than 10% of the top 100,000 websites, and which forms part of Google's Federated Learning of Cohorts (FLOC), the search giant's proposed alternative to cookie-based tracking. It's also a central technology in adtech platforms in general, therefore reaching far more than the 10% of sites identified in the above-mentioned study.

Category	Feature Name	Spoofable	Spoofing Approach	Detectable by Targeted Websites
C1	1. User-agent + - * [25]	Yes	a, b, c	Hard
	2. WebGL information - * [19]		b	
	3. System time + - * [25]		a, b	
	4. Battery information [48]		a, b	
	5. Cookie enabled + - * [12]		a, b, c	
	6. WebRTC [24]		b	
	7. Password autofill [46]		b	
	8. Platform - * [15]		a, b	
	9. Language + - * [23]		a, b, c	
	10. Local storage + - * [49]		b	
	11. Resolution + - * [11]		a, b	
	12. Do Not Track - * [16]		a, b, c	
C2	1. HTML5 canvas fingerprinting - * [44]	Yes	b	Hard
	2. System performance [43]		b	
	3. Font detection [31]		b	
	4. Scroll wheel fingerprinting [47]		b	
	5. CSS feature detection [42]		b	
C3	1. Browser plugin fingerprinting + - * [30]	Yes	a, b	Hard
	2. Browser extension fingerprinting [36]		b	

Typical facets that can be extracted from a user's browser without the need for cookies.

Identifiers that can be extracted from a user visit (collected via JavaScript APIs and HTTP headers) into a clonable browser profile include language settings, operating system, browser versions and extensions, installed plugins, screen resolution, hardware, color depth, time zone, timestamps, installed fonts, canvas characteristics, user-agent string, HTTP request headers, IP address and device language settings, among others. Without access to many of these characteristics, a great deal of commonly-expected web functionality would not be possible.

Extracting Information Via Ad Network Responses

The authors note that advertising data about the victim is quite easy to expose by impersonating their captured browser profile, and can be usefully exploited:

> *'[If] the browser fingerprinting is employed for personalized and targeted ads, the web server, hosting a benign website, would push the same or similar ads to the attacker's browser like the ones that would have been pushed to the victim's browser because the web server considers the attacker's browser as the victim's browser. Based on the personalized ads (e.g., related to pregnancy products, medications and brands), the attacker can infer various sensitive information about the victim (e.g. gender, age group, health condition, interests, salary level, etc.), even build a personal behavioral profile of the victim.*
>
> *'Leakage of such personal and private information can raise a frightful privacy threat to the user.'*

Since browser fingerprints change over time, keeping the user coming back to the attack site will keep the cloned profile up-to-date, but the authors maintain that a one-time cloning can still enable surprisingly long-term effective attack periods.

User Authentication Spoofing

Getting an authentication system to eschew two-factor authentication is a boon to cyber-criminals. As the authors of the new paper note, many current authentication (2FA) frameworks use a 'recognized' inferred browser profile to associate the account with the user. If the site's authentication systems are satisfied that the user is attempting to log in on a device that was used at the last successful login, it may, for user convenience, not demand 2FA.

The authors observe that Oracle , InAuth and SecureAuth IdP all practice some form of this 'check skipping', based on a user's recorded browser profile.

Fraud Detection

Various security services use browser fingerprinting as a tool to determine the likelihood that a user is engaged in fraudulent activities. The researchers note that Seon and IPQualityScore are two such companies.

Thus, it's possible, through the proposed methodology, to either unjustly characterize the user as a fraud by using the 'shadow profile' to trigger the thresholds of such systems, or else use the stolen profile as a 'beard' for genuine attempts at fraud, deflecting forensic analysis of the profile away from the attacker and towards the victim.

Three Attack Surfaces

The paper proposes three ways that the Gummy Browser system might be used against a victim: *Acquire-Once-Spoof-Once* involves appropriating the victim's browser ID in support of a one-time attack, such as an attempt to gain access to a protected domain in the guise of the user. In this case, the 'age' of the ID is irrelevant, since the information is acted on quickly and without follow-up.

In a second approach, *Acquire-Once-Spoof-Frequently*, the attacker is seeking to develop a profile of the victim by observing how web servers respond to their profile (i.e. ad servers that deliver specific types of content on the assumption of a 'familiar' user that already has a browser profile associated with them).

Finally, *Acquire-Frequently-Spoof-Frequently* is a longer-term ploy designed to regularly update the victim's browser profile by having the victim repeat their visit to the innocuous exfiltration site (which may have been developed as a news site or blog, for instance). In this way the attacker can execute fraud detection spoofing over a longer period of time.

Extraction and Results

The spoofing methods used by Gummy Browsers comprise script injection, use of the browser's setting and debugging tools, and script modification.

The characteristics can be exfiltrated with or without JavaScript. For instance, user-agent headers (which identify the brand of browser, such as *Chrome*, *Firefox*, et al.), can be derived from HTTP headers, some of the most basic and non-blockable information that's necessary for functional web browsing.

In testing the Gummy Browser system against FPStalker and Panopticlick, the researchers achieved an average 'ownership' (of an appropriated browser profile) of more than 0.95 across three fingerprinting algorithms, effecting a workable clone of the captured ID.

The paper emphasizes the need for systems architects not to rely on browser profile characteristics as a security token, and implicitly criticizes some of the larger authentication frameworks that have adopted this practice, especially where it is used as a method of maintaining 'user friendliness' by obviating or deferring the use of two-factor authentication.

The authors conclude:

> *'The impact of Gummy Browsers can be devastating and lasting on the online security and privacy of the users, especially given that browser-fingerprinting is starting to get widely adopted in the real world. In light of this attack, our work raises the question of whether browser fingerprinting is safe to deploy on a large scale.'*

Writer on machine learning, domain specialist in human image synthesis. Former head of research content at Metaphysic.ai, until its dissolution into DNEG's Brahma.ai.
Portfolio site: martinanderson.ai
Contact: martin@martinanderson.ai