

Impersonation-as-a-Service: Characterizing the Emerging Criminal Infrastructure for User Impersonation at Scale

Michele Campobasso

m.campobasso@tue.nl

Eindhoven University of Technology
Eindhoven, Netherlands

Luca Allodi

l.allodi@tue.nl

Eindhoven University of Technology
Eindhoven, Netherlands

ABSTRACT

In this paper we provide evidence of an emerging criminal infrastructure enabling impersonation attacks at scale. *Impersonation-as-a-Service* (IMPaaS) allows attackers to systematically collect and enforce user profiles (consisting of user credentials, cookies, device and behavioural fingerprints, and other metadata) to circumvent risk-based authentication system and effectively bypass multi-factor authentication mechanisms. We present the IMPaaS model and evaluate its implementation by analysing the operation of a large, invite-only, Russian IMPaaS platform providing user profiles for more than 260'000 Internet users worldwide. Our findings suggest that the IMPaaS model is growing, and provides the mechanisms needed to systematically evade authentication controls across multiple platforms, while providing attackers with a reliable, up-to-date, and semi-automated environment enabling target selection and user impersonation against Internet users as scale.

KEYWORDS

user profiling; impersonation attacks; impersonation-as-a-service; threat modeling

1 INTRODUCTION

In recent years there has been a surge in criminal infrastructures supporting cyberattacks and cybercrime activities at large [2, 10, 20]. For example, *exploitation-as-a-service* and *pay-per-install* provide a set of attack technologies generally aimed at infecting systems or controlling bots that are then employed to launch, for example, DDoS attacks, or subsequent malware and phishing campaigns (e.g., to harvest credit card numbers or steal credentials). An important problem in any venture, let alone a criminal one, is the ability to *systematically* monetize the effort that goes into it [22]. In criminal enterprises, monetization is not necessarily an easy feat: whereas re-selling or giving access to infected systems to fellow criminals alleviates the problem for whom generates the infection (e.g., the *bot herder* [5, 26]), the problem of assigning a price to each bot remains [3]. Whereas the dynamics of demand and offer in the underground are likely to play a role in this setting (and remain an important open question to investigate in this domain), another key factor in determining the value of an infected system is the information it manages and/or processes; for example, access to the email account(s) of an Internet user may have a different value, to attackers, than access to a user profile with a server-stored credit card number (e.g., an e-commerce website). On the other hand, it is not yet clear how (and if) attackers can *systematically* employ those credentials to impersonate Internet users at large, particularly in the presence of multi-factor authentication systems whereby a

username and password alone are not sufficient to gain access to an Internet account.

Credential theft and re-selling in underground communities have been studied multiple times in the literature; for example, recent studies provide an in-depth view of what happens to credentials after they have been stolen [35], and their employment for final attacks [40]. Similarly, several studies investigate the attack vectors that allow attackers to obtain the credentials in the first place, ranging from (targeted) phishing and phishing kits, to malware infections at scale [8, 9, 35]. On the other hand, a systematic employment of the stolen credentials remains out of reach for most attackers: credentials stolen from the underground may be accessed by multiple criminals, effectively destroying their value for later accesses [22]; similarly, the effort required to monetize access to stolen or hijacked user accounts does not scale well with the number of available accounts [22, 23]. In particular, protection systems such as multi-factor and risk-based authentication systems severely limit the capabilities of attackers to effectively employ stolen credentials, requiring the employment to more sophisticated attack vectors than a simple credentials dump [40]. Risk-based authentication systems receive user authentication requests and are responsible to decide whether additional multi-factor authentication is required for that session, or if the provided (valid) password suffices to grant access to the user requesting it. The idea behind risk-based authentication is that, by 'measuring' certain characteristics of the user environment (i.e., its fingerprint [1]), the authenticating system can build a 'risk profile' associated to that request as a function of the distance between the current fingerprint and the profile associated to the requesting user. If the mismatch is too large, the risk-based authentication system will defer the decision to a multi-factor mechanism (e.g., requesting a code sent to a trusted device or account, such as a mobile phone or an email account); on the other hand, if no anomaly in the user profile is detected, the risk-based authentication system will – in most cases – grant access just with the password.

This mechanism is a significant obstacle to a successful impersonation attack, as the very high dimensionality of a user fingerprint makes it impossible, for an attacker, to *systematically* reproduce it for arbitrary users from scratch [1, 40]. A recent study by Thomas et al. [40] highlights how modern phishing kits [34] are equipped with fingerprinting modules that, together with the user credentials, obtain a measurement of the user's environment that can be re-used to circumvent risk-based systems. On the other hand, obtaining these user profiles require systematic efforts to phish targets, perhaps across different platforms, and may not provide reliable and stable measures of a user's fingerprint as the victim's interaction with the attacker website may not accurately reflect the victim's

interaction with the legitimate website (e.g., for behavioural fingerprinting [12, 38]). Overall, traditional attack strategies seem unsuitable to reliably obtain, update, and enforce user profiles.

In this paper we provide evidence of a new emerging criminal infrastructure for *Impersonation-as-a-Service*, that relies on custom malware and a marketplace platform to systematize the delivery of complete *user profiles* to attackers. A user profile on an IMPaaS service comes complete with stolen credentials for multiple platforms, the ability to either reproduce or re-generate a user fingerprint from the stolen data, and a software bundle to enforce the user profile during an authentication session. To study the presence of the IMPaaS model in the wild, we provide an in-depth analysis of a large criminal platform (ImpaaS.ru) providing, at the time of writing, more than 260'000 profiles of Internet users, globally. ImpaaS.ru is an emerging, invite-only, Russian IMPaaS platform currently operating in the underground. To evaluate the nature of IMPaaS operations, we dissect the process behind the *acquisition, selection, and enforcement* of stolen user profiles enabled by the IMPaaS model, and provide a detailed evaluation of the characteristics of ImpaaS.ru, its extension, the characteristics of the user profiles it provides to final attackers, and the relative effect of different user profile characteristics on its value.

Scope and contribution. The contribution of this paper is three-fold:

- we provide the first characterization of the IMPaaS model for the systematization of impersonation attacks at scale;
- we provide an evaluation of a large, invite-only, emergent Russian IMPaaS platform that automates the collection, provision and enforcement of user profiles collected worldwide;
- we provide insights on the relative effects of different user profile characteristics on the value of the user profile, and quantify these effects.

A detailed technical analysis of the malware for the user profile exfiltration and enforcement is out of the scope of the present paper.

This paper proceeds as follows: Section 2 set the background for impersonation attacks and their relation to existent countermeasures; Section 3 introduces the IMPaaS model for impersonation attacks at scale, and Section 4 describes the ImpaaS.ru marketplace implementing it, and our infiltration and data collection strategy. ImpaaS.ru operations are analysed in Section 5. Section 6 discusses our findings, and Section 7 concludes the paper.

2 BACKGROUND AND RELATED WORK

2.1 User impersonation attacks

With the rise of sophisticated web applications, much of a user's Internet activity happens by accessing a multitude of remote services, from banking to e-commerce and social network platforms, through the browser. Most of these services will have authentication mechanisms that are meant to grant access to the underlying service to the authorized user(s) only. From an attacker's perspective, user impersonation provides a large portfolio of additional attack opportunities, ranging from economic gain [2, 16] to more targeted scenarios such as targeted-phishing [24] and violent crimes [21].

Password-based authentication (PBA) is the most common (first) barrier attackers have to overcome to perform an impersonation attack. Whereas passwords have proven difficult to securely handle,

are prone to leaks and to off-line attacks [32, 43] and still present severe usability problems [37], they represent the most widespread means of authentication online [7, 8]. PBA requires users to create a non trivial secret, not to reuse it across several services and to memorize both the secret and where it has been used; nonetheless, several studies indicate that up to $\approx 90\%$ of users reuse passwords or small variations thereof across several services [14, 28].

Whereas this leaves room for password guessing attack, additional attack vectors (such as malware and phishing [9, 40]) can be used to obtain user passwords, regardless of their complexity. In general, hijacked accounts can allow adversaries to tap into social connections of victims to compromise additional accounts [18, 39], by creating targeted social-engineering attacks against their circle of trust or by spamming malicious content [36], liquidate financial assets [27], steal sensitive information with the aim of blackmailing users [9, 36] and sextortion [42]. Additionally, stolen user credentials are oftentimes made available to the cybercrime community through underground markets [35, 40]. These markets generally provide 'dumps' of stolen credentials obtained from data leaks from an affected platform, or as a result of an extensive phishing campaign targeting its users [40]; common target platforms include banking or trading websites, cryptocurrency services, pornographic websites, and other internet services. A recent estimation calculates that, between March 2016 and March 2017, 1.9 billion phished credentials has been sold through the underground markets [40].

2.2 Countermeasures to attacks against PBA

Multi-Factor Authentication. To mitigate the shortcomings of authentication mechanisms relying solely on passwords, web platforms have started adopting additional authentication measures such as Multi-Factor Authentication (MFA). MFA moved the authentication paradigm from (solely) something that the user *knows* (e.g. a password) to something the user *has* (e.g., a token) [15, 40]. This is achieved mainly with a combination of a pair of valid credentials and a One Time Passcode (OTP) received via some trusted component such as a mobile phone, email, or a hardware token [15]. Albeit possible attack scenarios exist where the attacker can obtain the information required for the authentication almost in real-time (stolen token generator, compromised email, SIM swap attacks [33], etc.), MFA dramatically increases the costs for an attacker, and is widely regarded as an effective countermeasure to password-based impersonation attacks [40]. Nonetheless, MFA is not devoid of security problems, perhaps most notably related to its usability [31], concerns on token-recovery mechanisms, and third-party trust [7].

Risk-Based Authentication. Partly to mitigate the usability problem, *Risk-Based Authentication* (RBA) is oftentimes adopted as a means to evaluate whether the authenticating user is (likely to be) the one that has, historically, access to a specific account. RBA is an adaptive security technique aiming to strengthen password-based authentication by monitoring how unexpected or suspicious a login attempt is from the perspective of the authenticating service [31, 40, 41]. During the authentication, the RBA system monitors both behavioral and technical characteristics of the user and of the device, producing a *fingerprint* of the authenticating user [41]. RBA computes a risk score associated to the ongoing authentication by comparing the existent profile of the authenticating user

against the features collected for that instance of the authentication. The features vary from basic information such as User-Agent, system time and OS, to environmental or behavioral features, such as system language, keyboard layout, fonts and plugins installed, mouse movement, geolocation and keystroke speed [1, 17, 40, 41]. Whereas the high dimensionality of this data generates, with high probability, *unique* ‘fingerprints’ of a user, these are not necessarily *stable* in time (as, for example, users may access the service from multiple or new systems, may update software configurations, or authenticate from different locations). Depending on the computed risk score for that transaction, the authenticating service may grant access to the user with only a valid password (if the risk level is low), or require additional authentication factors (e.g., codes sent to associated email accounts, SMS verification) or even deny access for higher risk levels [31, 41]. This mechanism relies on the assumption that attackers cannot systematically re-create the profile of the victim, unless the attacker is already in control of a user’s system.

Following the implementation of RBA techniques across critical services, adversaries developed sophisticated solutions aiming to impersonate the user profile of the authenticating user. Recent literature has shown that phishing kits have developed capabilities to obtain user profiles that can then be re-used by the attacker; similarly, recent malware has been specifically engineered to report user activity back to the attacker [40]. In particular, Thomas et al. [40] highlight the improved capabilities of phishing kits in collecting information related to victims, including geographical location, browser metadata and answers to security questions; they found that attacks relying on user profile information collected from phishing kits are 40 times more likely to be successful than ‘regular’ attacks based on leaked credentials. On the other hand, the collection of user profile information does not scale well across users and platforms as user profiles may vary with time, across services, and must to be collected by the attacker through additional attack means (e.g., phishing).

2.3 Analysis of current attack strategies

Attack capabilities. From the analysis above, we identify six capabilities required to systematically bypass RBA systems.

Password authentication. At the very minimum, an attacker needs the authentication credentials of the victim.

User profiling. To attempt circumventing RBA systems, an attacker should have an accurate measurement of the victim’s profile/fingerprint for that platform.

Multi-platform. The attacker may need to access multiple web platforms to bypass some MFA controls (e.g. tokens or OTPs sent to an email account of the victim). Authentication credentials and user profiles need to be collected for these additional platforms as well. The capability of impersonating the victim on multiple platforms further increases the attack surface in scope of the attacker.

Profile updates. User profiles are unique but not necessarily stable. For example, a user may update a password, change software configuration, or access the service from a different geographical region. These changes may invalidate previously collected profiles for that user, which may therefore require updating.

Infection infrastructure. The attacker requires an infrastructure to infect users, and collect and update the collected user profiles. This has to be maintained as defensive capabilities evolve (e.g. blacklisting of an employed phishing domain), and may require the acquisition of external services (e.g., for an infection update [10, 20]).

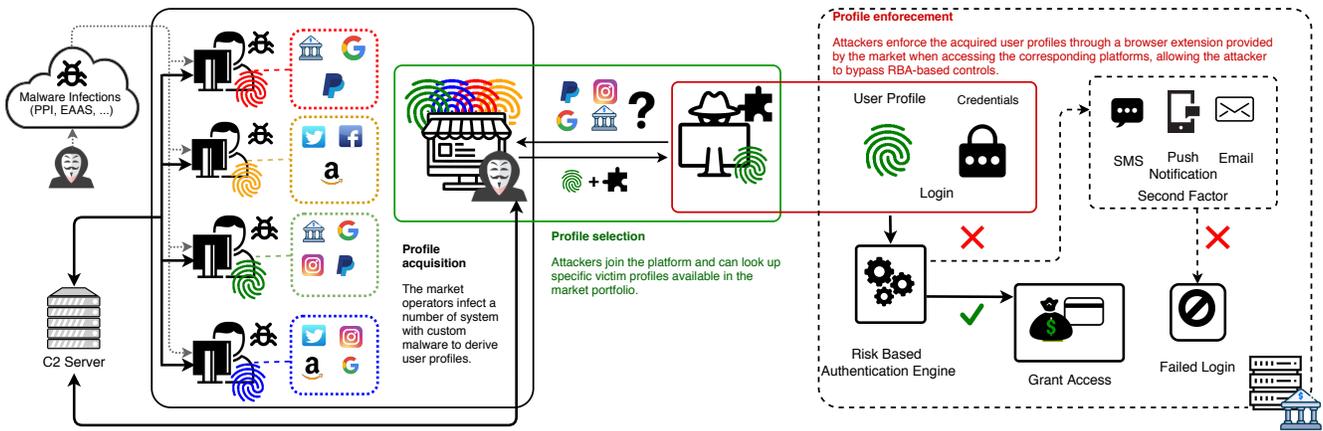
Automated profile enforcement. Once a profile is collected, the attacker needs to enforce it when authenticating on the platform. Whereas some aspects of the profile are easy to reproduce (e.g., user agent, screen resolution), others are not (e.g., installed fonts/plugins, keystroke speed, mouse movements, etc.). As profiles change across users and platforms, the attacker likely needs a system capable of enforcing the collected profiles in an automated fashion.

Analysis across attack strategies. Kurt et al. [40] identify three main strategies for impersonation attacks. Table 1 provides an overview of their capabilities.

Leaked credentials. credentials derived from data breaches on a platform. Leaked credentials are generally traded in bulk in underground forums; the leaked data oftentimes only contain associations between usernames and (hashed) password, with no user profile information. The data is static and if a user changes the password, the information owned by the attacker loses all value. As the leak concerns only one platform (and multiple leaks are likely unrelated to each other), cross-platform attacks against one user are not enabled by this attack strategy. However, password-reuse attacks may provide the attacker with access to additional platforms on top of the one that suffered the leak.

Phishing kits. attackers can employ kits to deploy phishing websites aimed at stealing user credentials. As users directly interact with the phishing kit, user profiling can be achieved by injecting fingerprinting code in the phishing webpage [40]. The profiles derived through phishing kits are however limited to only one occurrence of the authentication (on the phishing website) and may be incomplete or inaccurate. For example, the employment of password manager software may hinder the realism of the derived fingerprint (e.g., in terms of input time or user behaviour on the page) when compared to the one measured by the original platform. To achieve multi-platform capabilities, the attacker must develop or acquire a phishing kit for each of the phished platforms, and collect the relevant data through separate attacks against the same user.

Malware. the attacker has access to the system through a keylogger or trojan/bot. This requires the attacker to either purchase/rent the infected system [20], or create the infection themselves (e.g., through malware attached to a phishing email, or through Pay-per-Install services [10]). Due to the specificity of the attack, custom malware is likely needed to collect and update the profiles. As the attacker is virtually already in full control of the user system, they can collect user profiles related to any platform accessed by the victim. However, due to the position of the attacker, most of the impact (e.g., email access or web session hijacking) can be achieved through malware without the need of collecting the user profiles to then replicate them at a later stage.



By only owning the credentials of the victim, the attacker cannot bypass the MFA as the Risk-Based Authentication (RBA) system will detect an anomaly in the profile of the authenticating user. By relying on *Impersonation-as-a-Service* (IMPaaS), the attacker can reliably impersonate that profile by providing the values the RBA system expects for that user. IMPaaS obtains user profiles from a (large) botnet, and provides them in *bundles* as user profiles. An attacker purchases a user's profile(s) on the IMPaaS platform together with a browser extension that, provided the victim's profile as an input, reproduces it when accessing a service.

Figure 1: Diagram of Impersonation-as-a-Service operations.

Table 1: Overview of impersonation attack capabilities.

● indicates full systematic capability; ◐ indicates systematic capability only after specific engineering effort from attacker; ○ indicates no systematic capability.

	Leak	Phishing kits	Malware	ImpaaS
Password auth.	◐	●	●	●
User profiling	○	●	◐	●
Multi-platform	○	◐	◐	●
Profile updates	○	○	●	●
Infection infrastructure	○	○	◐	●
Automated profile enf.	○	○	○	●

3 THE IMPERSONATION-AS-A-SERVICE MODEL

In this paper we describe evidence of a new emerging attack model, namely *Impersonation-as-a-Service* (IMPaaS for short), and the criminal infrastructure supporting it.

IMPaaS directly addresses the main limitations of the ‘traditional’ impersonation attack strategies highlighted above by moving the acquisition and enforcement of victim profiles from an *ad-hoc* process to a *systematic* one. An overview of the comparison between IMPaaS and current vectors for impersonation attacks is summarized in Table 1. Figure 1 provides a birds’ eye view of the attack process, from *profile acquisition*, to *selection* and *enforcement*. IMPaaS operators rely on widespread malware infections to acquire ‘user profiles’ globally, and provide these profiles as ‘goods’ via the underground economy through a dedicated marketplace. As a result, attackers can acquire systematic access to a large set of user profiles spanning multiple platforms (social networks, email, corporate accounts, banking/cryptocurrency, etc.), alongside associated credentials and cookies; attackers can select the profiles they are most interested in based on a number of features, including the geographic location linked to the profile, the platforms for which impersonation data

is available, amount of stolen cookies, date of profile acquisition, and others. The user profiles available on the IMPaaS platform are automatically updated by the underlying infrastructure (e.g., as users change software configuration, or update passwords); further, the attacker can easily enforce and switch across the acquired user profiles by means of a dedicated browser extension provided by the IMPaaS platform, effectively commodifying the systematic impersonation of Internet users at large across multiple platforms.

Profile acquisition. The IMPaaS infrastructure is fueled by a botnet whose goal is, rather than solely collecting credit card information or banking credentials, to provide the information needed to replicate the user profiles of the infected victims across the online platforms on which affected users are active. The malware distribution is independent from the IMPaaS model: it can be delivered through phishing campaigns, targeted attacks, pay-per-install [10] or exploitation-as-a-service infrastructures [20]. Through the chosen attack vector, the attacker installs on the victim system custom malware engineered to collect user credentials and cookies from the victim’s browsers; the custom malware further collects a large set of technical and (user) behavioral information that can be replicated, by means of the infrastructure itself, to fully emulate the user; these include the fingerprint(s) of the victim’s browser(s) and other behavioral metadata that uniquely identify the user, such as network activity, browser history, cookie data, and interactions with the user interface of the platform. As profiles are fetched by means of a persistent malware infection, the infrastructure can provide updates of the profile data and credentials for each affected user. The harvested profiles and the respective updates are then pushed to the IMPaaS servers.

Profile selection. An IMPaaS operator provides the harvested user profiles to interested attackers via a dedicated marketplace. The marketplace provides an overview of the characteristics of the collected profiles available for purchase, such that the attacker can select which profiles best fit their goal by searching for victim profiles

that show specific features, such as a certain geographic location, web services for which stolen credentials are available, presence of cookies, etc. Albeit less targeted than allowed by a spear-phishing attack scenario, the selection procedure allows for a high degree of precision on the characteristics and/or environment of the user. For example, by browsing through the available credentials it is possible to identify users operating in a specific environment (e.g. a specific corporation, university, or other organizations), or with profiles on platforms of interest to the attacker. Once an attacker has identified their victim(s), the attacker can then proceed to buy the selected profiles. This can be achieved through the usual payment methods adopted in the cybercrime markets, such as via cryptocurrency payments to the marketplace, and/or by relaying the payment through a third-party escrow service. Importantly, as each profile can be purchased individually, the IMPaaS platform is in the position of removing purchased profiles from the marketplace listings, thus potentially reassuring the customer that they are the only one (next to the platform operators) with access to that profile.

Profile enforcement. The IMPaaS platform provides their customers with a customized software bundle that includes a custom browser (based on open-source projects) and a browser extension that allows attackers to fetch and ‘enforce’ the purchased user profiles during the attacker’s browsing session on that platform. Based on the profiles selected and purchased by the attacker, the software provided by the IMPaaS platform recreates a browsing environment that replicates the victim’s environment by instantiating exact copies of the stolen cookies and user credentials, and by spoofing other information on the victims’ systems (e.g., installed fonts/plugins, browser agent, ...). Further, the profile enforcement system provides cookies that embed behavioral metadata derived from the victim [12] without requiring explicit action from the attacker, and provides SOCKS5 proxy solutions to spoof the usual geographic location of the victim.

4 CHARACTERIZING IMPAAS IN THE WILD

In this section, we describe the operations of an emergent, invite-only IMPaaS platform, ImpaaS.ru¹. The platform has operated since late 2016 and grew considerably, in terms of available user profiles, in 2019. At the time of writing, ImpaaS.ru provides approximately 260’000 (and growing) user profiles available for impersonation attacks against Internet users worldwide. ImpaaS.ru is a Russian IMPaaS platform reachable from the surface web. This platform is, to the best of our knowledge, the first, large IMPaaS operator operating in the underground. On ImpaaS.ru, a user profile contains information coming from user systems infected with a credential stealer custom malware acting as a man-in-the-browser. The custom malware enables the exfiltration of cookies, credentials and sniffing of keystrokes, alongside additional environmental and device information that uniquely characterize the user. The IMPaaS platform states user profiles are updated and pushed to the attacker’s system in real-time, and that sold user profiles are removed from the listings of profiles available for purchase, although

¹We do not disclose the real name of the IMPaaS platform to minimize the risk of retaliatory actions from the market operators.

this is difficult to verify empirically, and ethically.² An overview of the profile characteristics is provided to browsing customers; profiles with specific characteristics can be searched through the marketplace interface. From the platform, it is possible to access the list of bought profiles and download the related fingerprint. Further, ImpaaS.ru provides their customers with a custom chromium-based browser plugin and a pre-built version of Chromium for both macOS, Linux and Windows. This bundle can be accessed only after having bought at least one user profile on the platform. The plugin comes with the capability of loading fingerprints previously obtained from the acquired profiles and can tunnel the traffic through an attacker-specified SOCKS5 proxy to spoof a victim’s geolocation.

Malware customization. The latest known custom malware employed by ImpaaS.ru is based on the AZORult malware [6, 13, 19]. ImpaaS.ru reports a recent update (Nov 2019) in AZORult addressing changes introduced in the Chrome browser that appear to have affected the malware functionality. Confirmation of massive phishing campaigns in that period associated with AZORult come independently from Kaspersky and other researchers [6, 19, 30]. Note that, start of 2020, AZORult was abandoned by ImpaaS.ru in favour of a new (and, at the time of writing, still unnamed), custom malware. Due to the changing nature of the adopted malware, we here only provide a high-level overview of AZORult operations from samples available (at the time of data collection) in the underground and malware repositories. For our analysis we replicated the latest three versions of AZORult (at the time of writing 3.3, 3.4.1 and 3.4.2) in a virtual environment, with the aim of evaluating its overall functionalities and their relevance to ImpaaS.ru. Malware customization happens through two modules, namely the *builder* and the *C2 server*. The *builder* has the purpose of generating the custom build of AZORult including the URL of the C2 server. The *C2 server* module is a ready-to-deploy web service providing an overview of the harvested data and a page for setting up the features of the malware; these features are user-defined, and include the collection of browser history, saved passwords, cryptocurrency client files, Skype history, a customizable regexp-based file grabber targeting user-defined folders on the infected host, and an additional setup for the deployment of a second stage infection on the victim system: as AZORult removes itself from the system after execution, the second-stage mechanism can allow ImpaaS.ru operators to obtain persistence on the infected system and further refine the data collection (e.g., to harvest behavioral data over time, see profile updates analysed in Sec 5).

4.1 Platform infiltration

Access to ImpaaS.ru is invite-only, and a valid account is needed to access the listings of available user profiles. Access to the registration procedure is provided through invite codes available to members already active on the platform, provided they spent at least 20 USD in purchased user profiles. To gain access to ImpaaS.ru we probed several underground forums in which we have a pre-existent foothold, and identified users that claim to be involved with ImpaaS.ru. As recent evidence suggests that underground

²A proposition is to infect one’s own system and purchase back the generated profile to verify its disappearance. As the malware employed by the platform is custom, reproducibility is non-trivial. See also Sec 4.

platform operators are actively monitoring and blacklisting ‘rogue’ accounts (e.g., performing scraping activities) [11], we aimed at the collection of several valid accounts prior to data collection to distribute the activity and have ‘backup’ identities to use if some of our accounts were to be blacklisted. Our search lead us to six members in Torum and one member in Crdclub (who claimed to be one of the operators of ImpaaS.ru) that were offering free invite codes between December 2019 and March 2020. We contacted them through the private messaging facility of the forums as well as on the messaging board, and obtained valid invitation codes from three of them in Torum. From Crdclub we gained access to an additional eight valid invitation codes using separate (and active) identities on the forum, for a total of eleven ImpaaS.ru accounts overall.

4.2 User profiles on ImpaaS.ru

ImpaaS.ru offers an overview of the available profiles, highlighting the information bundled in that user profile. A view of the interface accessed by attackers is provided in Figure 11 and Figure 10 in the Appendix. It is worth to note that, whereas ImpaaS.ru listings do not readily provide identifying information on the user, the information available on a listing is detailed enough to identify users operating in specific target environments such as a specific organization (e.g., to then perform lateral attacks [25]). ImpaaS.ru distinguishes between the following information in a user’s profile: cookies, resources and fingerprints.

Cookies. These are the cookies captured by the custom malware and available for injection toward the respective platforms once the user profile is purchased and enforced by the attacker.

Resources. Resources are collections of data derived from key-logging activity and probing of browser’s local resources, such as the database of stored passwords, and browser history. Some well-known resources (e.g., related to social media platforms, home banking, etc.) are highlighted as known resources by the platform, suggesting that the type of extracted Resources is an important information for the attacker to consider. A resource can include multiple data reporting login credentials, answers to security questions, detailed balance info for bank accounts, credit/debit card numbers and holder details. ImpaaS.ru states that the malware extracts Resources from infected systems through three main modules: FormParser reads the contents of the form data inputted by the user; SavedLogins gathers credentials saved in the browser’s local database; InjectScript implements code injection on the victim’s browser on behalf of the attacker, but its operation is unclear and most of the listed profiles do not appear to rely on it.

Fingerprints. Fingerprints provide a collection of the features exposed by a browser when interacting with RBA systems, ranging from technical metadata (user-agents, browser version) to more finely grained features (geolocation, latency, system language, fonts installed, web site device access permissions, etc.)³. Depending on the specific RBA implementation, a service may probe a specific subset of the features characterizing a browser or system. Differently

from Resources (which are tied to a specific service, e.g. a username/password combination on Amazon), the features collected in an ImpaaS.ru fingerprint are not bounded to a specific service, but to the browser environment itself (e.g., available system fonts, or installed plugins). Therefore, these constitute a *pool* of features that can be requested by any service, when available. ImpaaS.ru distinguishes between two types of Fingerprints:

- (1) **Real fingerprints:** these are directly collected from the victim’s device, providing an accurate identity of the impersonated device; albeit rarely available in bots, they appear to be sought after by market users;
- (2) **Synthetic fingerprints:** these fingerprints are generated on the basis of the data collected by the malware. However, accurate ‘synthetic’ fingerprints cannot be generated without user data (e.g. system fonts, plugins installed in a browser, etc.). For this reason, we consider the availability of Resources *and* of browser data in a user profile as an indication that the malware is in the position of collecting the necessary data to generate a reliable synthetic fingerprint.

4.3 Data collection strategy

To collect data on ImpaaS.ru operations, we first consider a number of structural limitations at the core to our sampling strategy:

- Lim-1 To avoid disclosing our identity to the ImpaaS.ru operators, we perform the scraping behind TOR. This poses technical limits (as well as ethical concerns) for bandwidth usage.
- Lim-2 We have a limited number of accounts to perform our measurements; aggressive probing risks exposing our accounts to the ImpaaS.ru operators, and lead to blacklisting.
- Lim-3 Information on Resources cannot be accessed in bulk via an API or other requests to ImpaaS.ru, but rather have to be requested in limited bundles with separate requests. This explodes the number of requests necessary to obtain Resources information on all user profiles on ImpaaS.ru.

To address Lim-1 and Lim-2, we employ an *ad-hoc* crawler. Initially the crawler was set to work $\approx 24\text{h/day}$ issuing, on average, 15 requests per minute; despite the relatively low requests volume, this strategy led two of our accounts to be blacklisted, suggesting that ImpaaS.ru operators may be employing network monitoring solutions to avoid measurement activities. Following [11], we progressively reduced the crawling activity to $\approx 6\text{h/day}$. In the process, an additional three accounts were banned, for a total of five banned accounts. It is interesting to note that three of the five blocked accounts were not linked to each other in any way,⁴ suggesting that market operators have kept their crawling-detection efforts high during our activities. To mitigate this problem we employed different strategies to access specific pages and resources to crawl on ImpaaS.ru: as already noted in [11], accessing URLs directly (as opposed to via website navigation) may generate anomalies in crawler monitoring systems. For this reason, we operationalised all crawling activities through browser instrumentation, and configured the crawler to mimic activity patterns compatible with those

³Whereas a full list of the probed features is not available from ImpaaS.ru nor from our analysis (see section 4), a number of commercial and free solutions could be employed by the ImpaaS.ru malware to implement reliable fingerprinting of the infected systems.

⁴The first two accounts were obtained from a single member of ImpaaS.ru and active on Torum. The other accounts all came through invitation codes generated by either different market members, or released by ImpaaS.ru operators themselves to different and unrelated accounts we control on Crdclub.

Table 2: Categories of resources.

Category	Definition	Examples
Services	Platforms providing the delivery of physical (e.g., goods, postage, etc.) or digital (e.g., content streaming, cloud, mail, etc.) services to final users.	Google, Postelitaliane
Social	Platforms to share user generated content.	Twitter, Skype
Money transfer	Platforms enabling direct payments between people or organizations using traditional payment circuits.	CreditUnion, Transfergo
Cryptoc.	Platforms enabling direct payments between people or organizations using cryptocurrency circuits.	Coinbase, Bittrex
Commerce	Platforms whose sole purpose is to purchase or book goods/services from one or multiple vendors.	Amazon, SaldiPrivati
Other	Platforms that do not match any of the previous categories.	Auth0

of a human user (e.g. timeouts between requests proportional to the length of the visited webpage, taking breaks, ...). With this final setup, we finally managed to silently crawl the market avoiding the detection and ban of the remaining accounts in our possession.

While necessary, the above strategy makes it impossible to gather complete information on Resources due to the exploding number of requests (Lim-3). This results in two datasets:

- **Full** database includes information on approximately 262'000 user profiles on ImpaaS.ru, including (infection, update) dates, prices, number of browsers for which resources are available, number of collected fingerprints for that user profile, and number of stolen cookies.
- **Sampled** database adds Resources information to a random selection of approximately 5% ($n = 13'512$) of the user profiles available on the market.⁵

The collected data is available for sharing to the research community at <https://security1.win.tue.nl>.

4.3.1 Analysis procedure. The data analysis in Section 5 is split in two subsections: in Sec. 5.1 we provide an overview of the data collected in the **Full** dataset, and characterize ImpaaS.ru operations by looking at its evolution, victim profile characteristics, profile updates, and pricing; in Sec. 5.2 we analyse the distribution and effect of Resources on pricing, as reported in the **Sampled** database. Standard sanity checks (e.g. on the regression results presented in Sec 5.2) are performed on all analyses. Reported logarithms are natural logarithms unless otherwise specified.

Manual resources classification. To factorize the type of resources reported in **Sampled** database in the analysis, we provide a classification of each resource in one of six categories. Table 2 lists the employed categories and their corresponding definitions. The classification was done manually by one of the authors over 454

⁵This fraction was originally set to 10%, however approximately half of the selected profiles were removed from the market during the data collection process.

Table 3: Summary statistics of the collected datasets.

	Variable	min	mean	max	sd
Full dataset ($n : 262'080$)	No browsers	0	1.58	10	1.02
	No cookies	0	1719.56	125198	1773.57
	No <i>real</i> fprnts	0	0.06	17	0.32
	Date infection [†]	12-12-17	20-11-19	16-03-20	157.81
	Date updated [†]	01-02-18	23-11-19	15-09-19	156.69
	Country	<i>char</i>	<i>char</i>	<i>char</i>	<i>char</i>
Random sample (Dec'17-Mar'20, $n : 13'512$)	Price (USD)	0.7	7.83	96	7.62
	No browsers	0	1.57	8	1
	No cookies	0	1782.01	26981	1735.24
	No <i>real</i> fprnts	0	0.12	9	0.54
	Date infection [†]	28-03-18	08-01-20	16-03-20	40.04
	Date updated [†]	12-11-19	14-01-20	13-06-20	37.62
	Country	<i>char</i>	<i>char</i>	<i>char</i>	<i>char</i>
	Price (USD)	0.7	8.84	63	8.17
	No resources	0	31.13	1322	46.63
	Crypto	0	0.07	18	0.6
Money Tr.	0	1.66	385	6.23	
Social	0	7.95	1322	18.73	
Services	0	16.64	560	24.44	
Commerce	0	4.66	296	11.12	
Other	0	0.15	16	0.88	

[†]: dates are reported in dd-mm-yy format. *sd* in days.

unique platforms for which Resources are reported in the dataset. The other author independently classified a random sample of 100 platforms, reaching an agreement score of 89%; after review, conflicts were resolved and the classification was updated accordingly. Additional random checks did not reveal any remaining mismatch.

Ethical considerations and limitations. No personally identifiable information is reported in our dataset. IP addresses of victims are masked on the platform, and no detailed information about the victims is available without purchasing a user profile. For obvious ethical concerns, we did not purchase any. Whereas this limits our analysis in that we do not have access to the software bundle provided by ImpaaS.ru, and cannot ascertain in detail the quality or operative aspects of the IMPaaS service provided by ImpaaS.ru, we are in the position of providing a full evaluation of the data is available to the attacker when browsing for victims.

5 DATA ANALYSIS

Table 3 provides an overview of the collected datasets.

Full dataset. The data collection spans from Dec 2017 to March 2020, involving approximately 262'000 user profiles. Most user profiles available on ImpaaS.ru target only one browser, with the top 5% targeting three browsers. Only 35 user profiles report data for more than six browsers in our data. Cookie distribution is similarly skewed. Profiles are distributed globally across 213 countries⁶, and prices range from 0.7 to 96 USD; 50% of the profiles cost at most 5 USD, whereas the priciest 5% are priced above 20 USD.

⁶Although there are only 195 recognized countries worldwide, ImpaaS.ru reports ISO 3166-1 codes, which do not distinguish sovereign nations from dependent territories.

Sampled dataset. This dataset reports data on 5.2% of the profiles available on ImpaaS.ru spanning from March 2018 till March 2020 ($n = 13'512$). For this dataset, we collected detailed information regarding the available resources. As this is a random sample, values are distributed similarly to Full dataset. Additionally, we extract information on the number and type of resources available for each profile. The average profile has upwards of 30 resources; most resources are of type Services, whereas Social and Commerce are less common. Cryptocurrency and Money transfer resources appear to be the least numerous in a profile.

5.1 Overview of ImpaaS.ru operations

To provide an overview of the IMPaaS operations conducted in the market we first look at the full dataset summarized at the top of Table 3. Interestingly, we find that approximately 12% of all profiles are *not* associated to a browser on the victim's system.⁷ As these profiles do not allow for impersonation attacks under the IMPaaS model, we remove those from further analysis. Relative to the number of user profiles, the number of available *real* fingerprints is surprisingly low, with only 4.3% of the available profiles having at least one. Note however that this refers only to Real fingerprints collected by the malware, not the Synthetic fingerprints that can be synthesized from user data (ref. Sec.4.2). Nonetheless, this suggests that (real) fingerprints, available browsers, cookies, and resources could be the driving force behind ImpaaS.ru activities.

Figure 2 provides an overview of the geographic distribution of the user profiles available on ImpaaS.ru and their median price per country. Most of the profiles belong to users in the United States of America and Europe, with a high fraction of EU countries showing volumes similar to those of the US. Users in Asian and African countries are comparatively less affected. As commonly seen in Russian cybercrime markets [2], ImpaaS.ru does not provide profiles for users in Russia, Ukraine, Belarus, and Kazakhstan (CIS countries). Furthermore, with the exception of Chad and the Central African Republic, the CIS countries appear to be some of the only unaffected countries, globally. Overall, median prices appear to vary from country to country rather than at a macro-regional level. For example, EU median prices seem to be higher in Spain ($m = 9.55, sd = 9.07$) and GB ($8.3, sd = 7.5$) than in Germany ($m = 7.21, sd = 8.21$) or Finland ($m = 6.96, sd = 6.68$). A set of Wilcoxon Rank-Sum tests evaluating the alternative hypothesis that SP and GB profile prices are higher than DE and FI ones confirms this observation ($p < 0.0001$). The high median price in Mauritania (26 USD) is caused by only one profile (with no fingerprints, two browsers, and four thousand cookies) available for that country.

The rate of appearance of new and updated user profiles on ImpaaS.ru is depicted in Figure 3. A clear upwards trends in terms of number of available profiles is visible, with a large jump in available profiles in November 2019 (coinciding with the 2019 spike in phishing campaigns distributing AZORult [6, 19, 30]). Overall, in Figure 3 we observe a sustained rate of new (black bar) and updated (orange bar) profiles, suggesting that the platform is systematically updating existing profiles, while adding new ones to the platform portfolio. We further investigate the time passing between time of

infection and (last) profile updates; Figure 4 shows the boxplot distribution of time passed between the infection and the last update received by the platform, plotted against the date of installation; in red, the upper bound of the maximum possible time in between. Overall the distribution appears relatively stable, with a median update time ranging between ten hours and four days. Unsurprisingly, recently acquired profiles are updated only after a few hours from acquisition; overall, the distribution suggests that profiles are kept updated on average over an extended period of time, ranging from a few days, to several months at the extreme of the distribution.

5.1.1 Analysis of profile values. Figure 5 reports the moving average of user profile prices as a function of time. The value of the traded profiles steadily increases as time passes, a signal of growth of the platform. In particular, profile prices seem to have doubled since November 2019, perhaps as an effect of the updated malware released in that period discussed at the start of this Section. Figure 6 reports the relation between the number of available Fingerprints in a profile and its price. The effect of an increased number of available fingerprints is, albeit positive, very limited. The average price seems to stabilize around the median value of 5 USD regardless of the number of fingerprints available in the profile, suggesting once again that other variables could be at play. We find no correlation between number of available browser and number of cookies and prices. This is not surprising, as these dimensions express little in terms of *which* identities of the victim the attacker may affect.

To further look at factors that may determine the value of a profile, we look at the impact of the geographic location to which the profile is linked. To do so, we investigate the relation between (log-transformed) profile prices and the wealth of the country in which the profile is located, expressed in terms of (log) GDP per capita (as reported by the World Development Indicators [4]). The intuition is that, the more 'valuable' a target is perceived to be, the greater the value of the corresponding profiles might be. Figure 7 reports the analysis. A positive and statistically significant correlation emerges, suggesting that profile prices are indeed correlated to the wealth of the respective country, perhaps a sign of the perceived value of that user profile ($corr = 0.4, p < 0.001$).

We note that some user profiles on ImpaaS.ru appear to be discounted at a rate of 30%. We do not find a clear-cut effect explaining which profiles are likely to be discounted (Table 6 in the Appendix).

5.2 The impact of Resources on profile pricing

We first look at the distribution of resource types in the Sampled dataset. As for the Full dataset, we remove from further analysis profiles that aren't associated to at least a browser of the victim's system and, in addition, profiles that don't contain any stolen resource, limiting the size of the dataset to $n = 12'052$. Table 4 provides an overview of the distribution of user profiles per category. Note that a profile can have resources that belong to more than one category. Overall, Services is the most commonly available resource type across user profiles. Resources in the Social and Commerce categories are also common, with respectively about 70% and 40% of user profiles with resources in these categories. Approximately 25% of the profiles have data for banking and payment accounts; by contrast, less than 2% of user profiles have resources in the Cryptocurrency category. Only 4.5% of the resources in our

⁷Note that all profiles without browser data also do not, by definition, report any data on cookies or fingerprints.

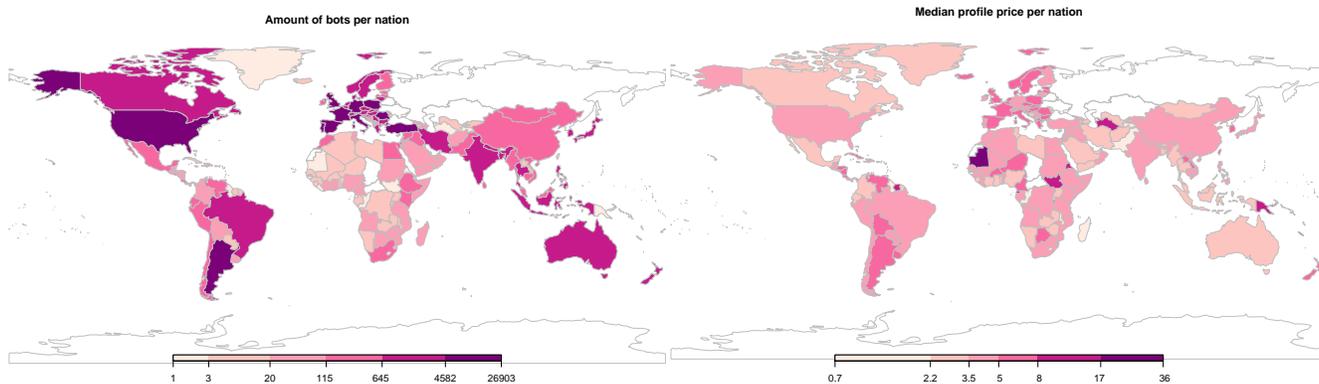


Figure 2: Global distribution of user profiles (left) and their median price (right) on ImpaaS.ru

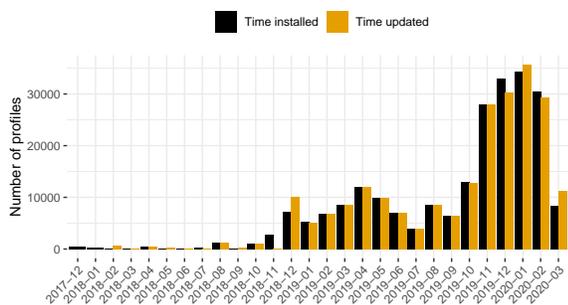


Figure 3: Progression of available user profiles over time.



Figure 5: Weekly moving average of user profile prices.

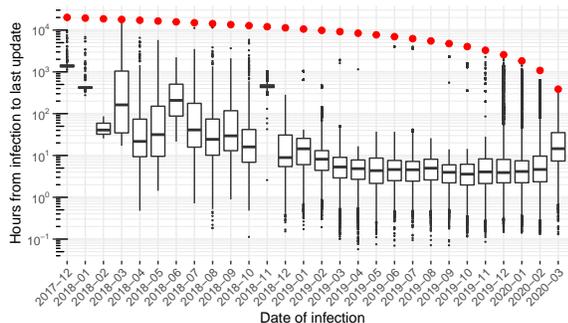


Figure 4: Time between infection and last profile update (in log scale).

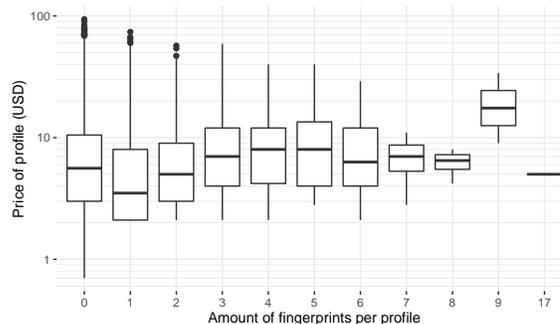


Figure 6: Relation between the amount of fingerprints available and the average price of user profile (in log scale).

dataset were classified as Other, indicating that the proposed classification covers the vast majority of resource types in ImpaaS.ru.

Figure 8 provides a first overview of the relation between the number of resources available in a profile and the associated price. A clear correlation emerges. The depicted linear log-log relation indicates negative marginal returns for each added resource, meaning that every additional resource added to a profile provide an increasingly smaller, albeit positive, added value to the profile. Further exploring the impact of resources on pricing, Figure 9 shows

the impact of the presence of resources in any specific category on the value of a user profile. Note that, because each profile can contain resources of more than one category, one cannot isolate the relative importance of each category here. However, the comparison shows how, on average, a profile that contains (also) resources in that category is priced versus other profiles that do not have it. This is meaningful as the categories show relatively low correlations (reported in Table 7 in the appendix). On average, profiles

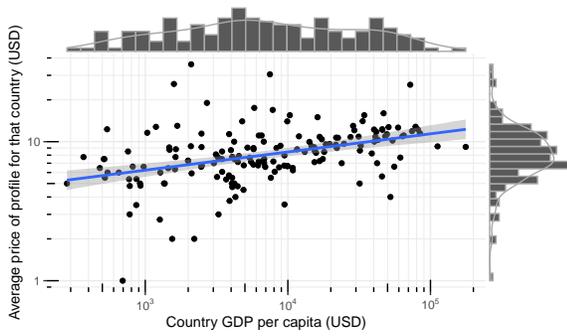


Figure 7: Relation between GDP per capita and average price of user profiles in that country (in log scale).

Table 4: Type of resources per user profile.

Resource type	no. profiles ($n = 12'052$)
Cryptocurrency	236
Money Transfer	3109
Commerce	5'066
Social	8'111
Services	11'167
Other	548

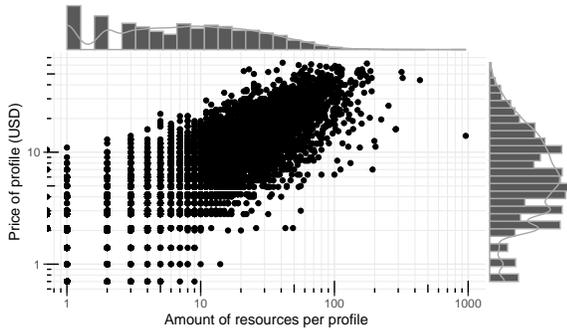


Figure 8: Relation between the amount of resources available and the average price of user profile (in log scale).

that include Cryptocurrency resources seem to be the most valuable. Money transfer and Commerce resources belong to profiles of approximately the same value, whereas profiles with Social and Services are the least valued in ImpaaS.ru. By comparing the relative ‘jump’ introduced by the addition of each category, one can further evaluate the added value, on average, of having a resource of that type. In this respect, Other appears to be the least ‘impactful’ category, as the appearance of a resource of this type is related to the smallest relative increase in price, on average, in a profile. On the contrary, Cryptocurrency and Money transfer resources cause the highest jump in profile value, passing from a median value of approximately 7 USD to more than 20. Other categories show less extreme changes in price. Overall, we find that resources associated to financial platforms and services appear to

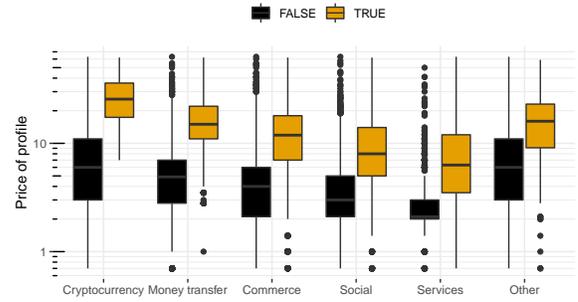


Figure 9: Profile price variation according to the presence of not of resources of a certain category (in log scale).

have the highest impact on the value of a profile, with Social and Services being the least valued. On the other hand, the addition of resources of any category appears to have a positive impact on the value of a profile.

To formally evaluate this relation, we build a set of linear regression models to quantify the effect of different profile features on profile values in ImpaaS.ru. To evaluate the effect of each factor independently, and monitor its relation to other characteristics of a user profile, we define the following nested models with response variable $y = price$ (the error term ϵ_i is omitted for brevity):

$$M1 : y_i = \beta_0 + \beta_1 \text{Real fingerprints}_i$$

$$M2 : y_i = \dots + \beta_2 \log(\text{GDP}_i)$$

$$M3 : y_i = \dots + \beta_3 \text{Cryptocurrency}_i + \beta_4 \text{Money transfer}_i + \beta_5 \text{Commerce}_i + \beta_6 \text{Social}_i + \beta_7 \text{Services}_i + \beta_8 \text{Other}_i$$

$$M4 : y_i = \dots + \beta_9 \text{resources}_i$$

where β_0 is the intercept, Real fingerprints is the number of fingerprints embedded in that user profile, $\log(\text{GDP})$ is the natural logarithm of the gross domestic product per capita for the country associated with the user profile, {Cryptocurrency...Other} are dummy variables representing the presence of resources of the corresponding category, and resources is the overall number of resources in that profile (irrespective of category).

Regression results are summarized in Table 5. To evaluate the effects of profile characteristics on full prices we remove profiles ‘on sale’ from the dataset. Table 8 and Table 9 in the appendix report, respectively, a full breakdown of the variables’ impact on the prediction, and the regression results for all data points including profiles on ‘sale’; both tables report results quantitatively and qualitatively in line with those reported in Table 5. Overall, the coefficient estimates appear stable across the models, with the exception of β_2 ($\log(\text{GDP})$), that becomes less important on the estimation of the dependent variable price as the types of Resources are added to the model. The change ranges from an expected increase of 0.2 USD in profile value for every 10% increase in GDP ($\beta_2 = 2.29$ in M2, $2.29 \times \log(1.10) = 0.22$), to a relatively smaller (0.04 USD) price increase when all resource categories are added in the model. This indicates that some resource categories may appear more frequently for high GDP countries than for others; with reference to Table 8

Table 5: Regression analysis on prices of user profiles.

	Model 1	Model 2	Model 3	Model 4
β_0	10.41*** (0.11)	-12.11*** (1.21)	-5.57*** (0.81)	-3.70*** (0.63)
Real Fngrpr	0.55*** (0.16)	0.69*** (0.16)	1.31*** (0.10)	1.11*** (0.07)
log(GDP)		2.29*** (0.12)	0.46*** (0.08)	0.42*** (0.06)
Crypto			13.62*** (0.44)	10.12*** (0.34)
Money Transfer			8.86*** (0.16)	6.20*** (0.13)
Commerce			5.06*** (0.15)	3.22*** (0.12)
Social			3.44*** (0.15)	1.68*** (0.12)
Services			3.95*** (0.29)	2.31*** (0.22)
Other			4.22*** (0.31)	0.89*** (0.24)
Resources				0.10*** (0.00)
R ²	<0.01	0.05	0.65	0.79
Adj. R ²	<0.01	0.05	0.65	0.79
Num. obs.	7123	7123	7123	7123

*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

in the appendix, it appears that resources of type Money transfer and Commerce tend to appear more often in wealthy countries, as most of the effect of the GDP variable disappears when this category is accounted for, while the opposite effect emerges when Social resources are included in the model. Additional resource categories have modest effects on the GDP coefficient estimate. As resource categories are added to the model, the impact of the number of fingerprints increases, passing from a 0.55 USD increase in expected profile value for each additional fingerprint in the profile ($\beta_1 = 0.55$) to a 1.31 USD increase estimated by M3. This suggests a positive joint effect of the number of fingerprints in a profile, and the number of platforms with resources an attacker can employ to impersonate a victim. All resources have a positive effect on the value of a user profile with Cryptocurrency and Money transfer having the highest impact, increasing the expected value of 13.62 and 8.86 USD respectively when available. Following this trend, Commerce shows a relatively large effect as well, increasing the profiles' expected value of 5.06 USD. These findings may not come as a surprise, and may indicate that ImpaaS.ru customers may be primarily aiming at economical profit (supporting insights from observing ImpaaS.ru customers discussing on a dedicated Telegram channel, see Sec. 6.2 for an informal report). Finally, the effect of the number of resources in M4 is significant and positive; interestingly, its addition decreases the effect of the single resource categories,

confirming the intuition that the more platforms an attacker can impersonate, the higher the value of the profile.⁸

In all, resource types appear to explain the majority of the variance in the model, with Money transfer accounting for a jump in more than 30% in the model (adjusted) R^2 when compared with previous model. The complete model explains most of the price variance in our dataset ($R^2 = 0.79$), suggesting that the model provides an appropriate description of the features determining user profile values in ImpaaS.ru.

6 DISCUSSION

In this paper we presented the IMPaaS model as a novel threat enabling attackers to perform user impersonation at scale. IMPaaS is supported by an emergent criminal infrastructure that controls the supply chain of user profiles, from system infection to profile acquisition and commodification. Whereas traditional impersonation attacks relying solely on stolen credentials are greatly mitigated by risk-based and two-factor authentication systems, the capability of seamlessly reproducing a user's 'appearance' to an authentication system allows attackers to systematically compromise accounts of multiple users, across multiple platforms.

Whereas Thomas et al. already suggested that user profiling could be used to bypass modern authentication systems [40], in this work we provide evidence of an emergent *as-a-service* impersonation model that appears to be rapidly expanding. The profile value analysis provided in Section 5.2 suggests a mature pricing model, which may indicate that the analyzed platform operations are of stable, predictable quality, and likely to expand in number. Overall, the analysis of the available user profiles on ImpaaS.ru and the reportedly widespread adoption of *info-stealer* malware such as AZORult in phishing campaigns [19, 29] provide further supporting evidence of the growth of this threat model.

Our analysis of ImpaaS.ru allows us to further quantify the relative effects of different resources on the value of a user profile. Interestingly, albeit perhaps not surprisingly, we find that profile values show a significant correlation with the wealth of the country (expressed in terms of GDP) associated to that profile; this suggests that attackers looking to impersonate and, likely, monetize user profiles assign a greater value to profiles likely to give access to greater financial resources (e.g., bank balances or valid credit cards). Interestingly, this effect is significantly reduced by the presence of Commerce resources in a profile, perhaps due to the prevalence of e-commerce platforms in wealthy countries. Nonetheless, other resource categories have a clear impact on the overall valuation of a user profile, with Cryptocurrency and Money transfer resources driving most of the value. *Real* fingerprints (those derived directly from the device, rather than being synthesized by the IMPaaS platform using the profile's metadata) available in a profile also add value to the user profile. Our analysis suggests that each *real* fingerprints adds about 0.55 USD to the value of a user profile, and up to 1.31 USD when considered jointly with the available resources,

⁸Driven by observations in Chen et al. [12], who identified cookies as having a key role in behavioral fingerprinting practices, we find that in terms of profile pricing the availability of cookies does not show a statistically significant effect (Anova $F_{1,92}$, $p = 0.17$) in our dataset, suggesting that cookies do not play a central role in impersonation attacks as driven by ImpaaS.ru.

suggesting that the *modus operandi* enabled by IMPaaS described in Figure 1 is supported by the platform operations.

Importantly, our analysis allows us to put a number on the value of user information to attackers, contributing to the literature on the subject. A user’s ‘virtual identity’ seems to be worth between less than 1 USD and approximately 100 USD. This value changes significantly depending on the wealth of the country where the user (appears to be) located; a rule-of-thumb indication seems to be that for a tenfold increase of the a user’s ‘expected wealth’ (approximated by a country’s GDP), a profile value increases on average by approximately 1 USD. Cybercriminals seem to particularly value profiles with access to Cryptocurrency and Money transfer platforms, whose prices are respectively 10 and 6 USD higher than profiles with no access to platforms of these types. To put this in perspective, these represent respectively a 150% and 90% markup over the price of the average profile, a clear sign of the relevance of resources of these types for cybercriminal activities. By contrast, access to Social and other services does not seem to be (in comparison) highly valued by cybercriminals.

6.1 Implications for victimization

The systematization of impersonation attacks enabled by the IMPaaS model allows attackers to select and target specific victim profiles, and to automate the attack procedure by means of dedicated software bundles replicating a victim’s browsing conditions on the environment of the attacker. Differently from traditional phishing-based attacks, IMPaaS provides an attacker with access to several platforms on which a user is active, effectively allowing the attacker to both mitigate security measures (e.g., by monitoring email for authentication codes or activity notifications), as well as extending the attack surface to different services (e.g. banking, social, etc.).

Attackers leveraging an IMPaaS platform can rely on an automated source for credentials to conduct sophisticated attacks at scale. In addition to obtaining access to banking websites, cryptocurrency exchange platforms and e-commerce websites, an attacker may compromise multiple accounts to gain control over the identity of the victim. The capability of selecting victim characteristics before the acquisition of a profile is also a potential enabler of targeted attacks against organizations or communities for which a victim is an employee, or a registered member. The attacker may employ that advantage point to facilitate lateral movement attacks, for example targeting colleagues or family members of the victim by using their legitimate contact details. Furthermore, the attacker could integrate additional information about a victim gathered through the accessed platforms (part of a corporation, subscription to meeting websites, etc.) to further escalate the attack to other victims.

6.2 Examples of (alleged) criminal operations enabled by ImpaaS.ru

To informally investigate how attackers are weaponizing capabilities enabled by IMPaaS, we collected a number of examples provided by users of ImpaaS.ru on a Telegram channel linked to the platform, to which we have gained access through ImpaaS.ru. Many attacks reported there appear to focus on Money transfer and Commerce services. For example, a user shared that they were

(allegedly) able to cash-out from a US bank using a synthetic fingerprint acquired on ImpaaS.ru and with the support of a geographically accurate SOCKS5 proxy. The user further suggested to rely on 911.re as the marketplace where to buy proxies linked to specific ZIP codes and/or ISPs. On a similar line, a second user reported to have managed issuing a new debit card on behalf of the victim, with the aim of cashing it out through ATMs. Interestingly, some ImpaaS.ru users report performing multi-stage attacks deployed through the obtained user profiles and exploiting multiple platforms. For example, an attacker describes setting filters to a victim’s email mailboxes accessed through the victim’s user profile, with the aim of hiding notifications from Amazon related to purchases the attacker made using the victim’s Amazon account.

Overall, whereas of course none of these examples can be verified and the threats described above are not new per-se, the mix of infrastructural support for profile acquisition, selection, and enforcement enabled by IMPaaS opens to the *systematization* of threat scenarios such as the ones described above, on a global scale.

7 CONCLUSION

In this paper we presented the emergence of the *Impersonation-as-a-Service* criminal infrastructure, which provides user impersonation capabilities for attackers at large. IMPaaS allows attackers to bypass risk-based authentication systems by automatically simulating the victim’s environment on the attacker’s system. In this study we characterise the largest currently operating IMPaaS infrastructure, ImpaaS.ru, by performing an extensive data collection spanning more than 260 thousand stolen user profiles collected worldwide. ImpaaS.ru infiltration and data collection required substantial efforts to collect multiple accounts, needed to fine-tune the data collection as platform operators seemed to monitor crawling activities and blacklist related accounts. From our analysis, ImpaaS.ru emerges as a mature, expanding infrastructure with a clear pricing structure, suggesting a well-established criminal business model. *Impersonation-as-a-Service* represents an additional component of the cybercrime economy, providing a systematic model to monetize stolen user credentials and profiles.

Lesson learned. Our data collection efforts provide supporting evidence that underground platform operators are actively monitoring crawling activities, and take measures to limit them. This may prevent future research activities and significantly impact the possibility of designing large-scale studies studying cybercriminal online venues. Specific sampling strategies and analysis techniques will have to be devised to further develop research in this domain.

ACKNOWLEDGEMENTS

This work is supported by the ITEA3 programme through the DEFRAUDify project funded by Rijksdienst voor Ondernemend Nederland (grant no. ITEA191010).

REFERENCES

- [1] ALACA, F., AND VAN OORSCHOT, P. C. Device fingerprinting for augmenting web authentication: classification and analysis of methods. In *Proceedings of the 32nd Annual Conference on Computer Security Applications* (2016), pp. 289–301.
- [2] ALLODI, L. Economic factors of vulnerability trade and exploitation. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security* (2017), ACM, pp. 1483–1499.

- [3] ANDERSON, R., AND MOORE, T. The economics of information security. *Science* 314 (2006).
- [4] BANK, T. W. World development indicators. <https://datacatalog.worldbank.org/dataset/world-development-indicators>.
- [5] BINSALLEH, H., ORMEROD, T., BOUKHTOUTA, A., SINHA, P., YOUSSEF, A., DEBBABI, M., AND WANG, L. On the analysis of the zeus botnet crimeware toolkit. In *Privacy Security and Trust (PST), 2010 Eighth Annual International Conference on* (2010), IEEE, pp. 31–38.
- [6] BISSON, D. Azorult trojan disguised itself as fake protonvpn installer, Feb 2020.
- [7] BONNEAU, J., HERLEY, C., VAN OORSCHOT, P. C., AND STAJANO, F. The quest to replace passwords: A framework for comparative evaluation of web authentication schemes. In *2012 IEEE Symposium on Security and Privacy* (2012), IEEE, pp. 553–567.
- [8] BONNEAU, J., HERLEY, C., VAN OORSCHOT, P. C., AND STAJANO, F. Passwords and the evolution of imperfect authentication. *Communications of the ACM* 58, 7 (2015), 78–87.
- [9] BURSSTEIN, E., BENKO, B., MARGOLIS, D., PIETRASZEK, T., ARCHER, A., AQUINO, A., PITSILLIDIS, A., AND SAVAGE, S. Handcrafted fraud and extortion: Manual account hijacking in the wild. In *Proceedings of the 2014 conference on internet measurement conference* (2014), pp. 347–358.
- [10] CABALLERO, J., GRIER, C., KREIBICH, C., AND PAXSON, V. Measuring pay-per-install: The commoditization of malware distribution. In *Usenix security symposium* (2011).
- [11] CAMPOBASSO, M., BURDA, P., AND ALLODI, L. Caronte: crawling adversarial resources over non-trusted, high-profile environments. In *2019 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)* (2019), IEEE, pp. 433–442.
- [12] CHEN, Y., PAVLOV, D., AND CANNY, J. F. Large-scale behavioral targeting. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining* (2009), pp. 209–218.
- [13] CYLANCE. Threat spotlight: Analyzing azorult infostealer malware, Jun 2019.
- [14] DAS, A., BONNEAU, J., CAESAR, M., BORISOV, N., AND WANG, X. The tangled web of password reuse. 2014. *Cited on* (2014), 7.
- [15] DMITRIENKO, A., LIEBCHEN, C., ROSSOW, C., AND SADEGHI, A.-R. On the (in)security of mobile two-factor authentication. In *Financial Cryptography and Data Security* (Berlin, Heidelberg, 2014), N. Christin and R. Safavi-Naini, Eds., Springer Berlin Heidelberg, pp. 365–383.
- [16] FRANKLIN, J., PAXSON, V., PERRIG, A., AND SAVAGE, S. An inquiry into the nature and causes of the wealth of internet miscreants. In *Proc. of CCS'07* (2007), pp. 375–388.
- [17] FREEMAN, D., JAIN, S., DÜRMUTH, M., BIGGIO, B., AND GIACINTO, G. Who are you? a statistical approach to measuring user authenticity. In *NDSS* (2016), pp. 1–15.
- [18] GAO, H., HU, J., WILSON, C., LI, Z., CHEN, Y., AND ZHAO, B. Y. Detecting and characterizing social spam campaigns. In *Proceedings of the 10th ACM SIGCOMM conference on Internet measurement* (2010), pp. 35–47.
- [19] GATLAN, S. Azorult malware infects victims via fake protonvpn installer, Feb 2020.
- [20] GRIER, C., BALLARD, L., CABALLERO, J., CHACHRA, N., DIETRICH, C. J., LEVCHENKO, K., MAVROMMATIS, P., MCCOY, D., NAPPA, A., PITSILLIDIS, A., PROVOS, N., RAFIQUE, M. Z., RAJAB, M. A., ROSSOW, C., THOMAS, K., PAXSON, V., SAVAGE, S., AND VOELKER, G. M. Manufacturing compromise: the emergence of exploit-as-a-service. In *Proc. of CCS'12* (2012), ACM, pp. 821–832.
- [21] HAVRON, S., FREED, D., CHATTERJEE, R., MCCOY, D., DELL, N., AND RISTENPART, T. Clinical computer security for victims of intimate partner violence. In *28th USENIX Security Symposium (USENIX Security 19)* (Santa Clara, CA, Aug. 2019), USENIX Association, pp. 105–122.
- [22] HERLEY, C. So long, and no thanks for the externalities: the rational rejection of security advice by users. In *Proc. of NSPW'09* (2009), NSPW '09, ACM, pp. 133–144.
- [23] HERLEY, C. Why do nigerian scammers say they are from nigeria? In *Proc. of WEIS'12* (2012).
- [24] HO, G., CIDON, A., GAVISH, L., SCHWEIGHAUSER, M., PAXSON, V., SAVAGE, S., VOELKER, G. M., AND WAGNER, D. Detecting and characterizing lateral phishing at scale. In *28th USENIX Security Symposium (USENIX Security 19)* (Santa Clara, CA, Aug. 2019), USENIX Association, pp. 1273–1290.
- [25] HO, G., JAVED, A. S. M., PAXSON, V., AND WAGNER, D. Detecting credential spearphishing attacks in enterprise settings. In *Proceedings of the 26rd USENIX Security Symposium (USENIX Security'17)* (2017), pp. 469–485.
- [26] HOLT, T. J., SMIRNOVA, O., AND HUTCHINGS, A. Examining signals of trust in criminal markets online. *Journal of Cybersecurity* (2016), tyw007.
- [27] IOACTIVE. Technical white paper: Reversal and analysis of zeus and spyeye banking trojans, 2012.
- [28] ION, I., REEDER, R., AND CONSOLVO, S. âÄIj... no one can hack my mindâÄI: Comparing expert and non-expert security practices. In *Eleventh Symposium On Usable Privacy and Security (SOUPS) 2015* (2015), pp. 327–346.
- [29] KREBS, B. Krebs on security, Mar 2020.
- [30] LABS, K. New azorult campaign abuses popular vpn service to steal cryptocurrency, Feb 2020.
- [31] MILKA, G. Anatomy of account takeover. In *Enigma 2018 (Enigma 2018)* (Santa Clara, CA, Jan. 2018), USENIX Association.
- [32] MORRIS, R., AND THOMPSON, K. Password security: A case history. *Communications of the ACM* 22, 11 (1979), 594–597.
- [33] MULLINER, C., BORGAONKAR, R., STEWIN, P., AND SEIFERT, J.-P. Sms-based one-time passwords: attacks and defense. In *International Conference on Detection of Intrusions and Malware, and Vulnerability Assessment* (2013), Springer, pp. 150–159.
- [34] OEST, A., SAFEI, Y., DOUPÉ, A., AHN, G.-J., WARDMAN, B., AND WARNER, G. Inside a phisher's mind: Understanding the anti-phishing ecosystem through phishing kit analysis. In *2018 APWG Symposium on Electronic Crime Research (eCrime)* (2018), IEEE, pp. 1–12.
- [35] ONAOLAPO, J., MARICONTI, E., AND STRINGHINI, G. What happens after you are pwnd: Understanding the use of leaked webmail credentials in the wild. In *Proceedings of the 2016 Internet Measurement Conference* (2016), pp. 65–79.
- [36] SABILLON, R., CAVALIER, V., CANO, J., AND SERRA-RUIZ, J. Cybercriminals, cyberattacks and cybercrime. In *2016 IEEE International Conference on Cybercrime and Computer Forensic (ICCCF)* (2016), IEEE, pp. 1–9.
- [37] STOBERT, E. The agony of passwords: Can we learn from user coping strategies? In *CHI'14 Extended Abstracts on Human Factors in Computing Systems*. ACM New York, NY, USA, 2014, pp. 975–980.
- [38] STRINGHINI, G., AND THONNARD, O. That ainâÄZt you: Blocking spearphishing through behavioral modelling. In *International Conference on Detection of Intrusions and Malware, and Vulnerability Assessment* (2015), Springer, pp. 78–97.
- [39] THOMAS, K., LI, F., GRIER, C., AND PAXSON, V. Consequences of connectivity: Characterizing account hijacking on twitter. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security* (2014), pp. 489–500.
- [40] THOMAS, K., LI, F., ZAND, A., BARRETT, J., RANIERI, J., INVERNIZZI, L., MARKOV, Y., COMANESCU, O., ERANTI, V., MOSCICKI, A., ET AL. Data breaches, phishing, or malware? understanding the risks of stolen credentials. In *Proceedings of the 2017 ACM SIGSAC conference on computer and communications security* (2017), pp. 1421–1434.
- [41] WIEFLING, S., IACONO, L. L., AND DÜRMUTH, M. Is this really you? an empirical study on risk-based authentication applied in the wild. In *IFIP International Conference on ICT Systems Security and Privacy Protection* (2019), Springer, pp. 134–148.
- [42] WITTES, B., POPLIN, C., JURECIC, Q., AND SPERA, C. Sextortion: Cybersecurity, teenagers, and remote sexual assault. *Center for Technology at Brookings* (2016).
- [43] YAN, Q., HAN, J., LI, Y., DENG, H., ET AL. On limitations of designing usable leakage-resilient password systems: Attacks, principles and usability. In *19th Network and Distributed System Security Symposium (NDSS)* (2012).

APPENDIX

A MARKET FEATURES

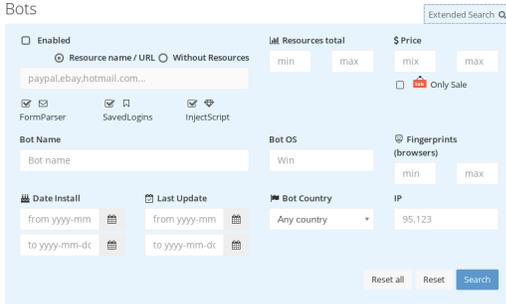


Figure 10: View on the advanced search functionality.

Figure 10 and Figure 11 report screenshots from the market. Figure 10 depicts the search function of the market. Attackers can access a fine grained research tool that enables them to search for profiles with specific resource composition, number of available browser fingerprints, and other information. In Figure 11, an overview of the details for each profile is provided. On the left, from top to bottom, the name of the profile and the installation and update dates of the profile are listed. On top, in the center, the list of the available browsers (here only Microsoft Edge). The superimposed number indicates the number of fingerprints available for that specific browser and the superimposed icon whether cookies are available (green) or not (red). On the top-right is reported the number of resources available (here 76). In the center, an overview of the websites for which resources are present. On the right, details about the country, IP prefix and operating system are provided. Finally, to the right-most, there’s the price expressed in USD and buttons to respectively buy, reserve or add the profile to the cart.

B FURTHER DATA INSIGHTS

Table 6 shows correlation coefficients for the logistic model $sale = \beta_0 + \beta_1 Resources + \beta_2 Year + \beta_3 Cookies + \beta_4 Browsers$ (binary response variable). Whereas Resources, Year, and Browsers are significant predictors, the effect is very small with the unsurprising exception of Year, suggesting that profiles recently acquired are *less* likely to be put on ‘sale’. The coefficient for β_4 shows a positive, albeit small, effect of the no. of browsers provided in a profile on the likelihood of profile sale. As indicated by the small R^2 , we do not find a clear rationale explaining this effect.

Table 7 reports correlation coefficients between Resources types in our dataset. No high correlation is found, suggesting that no auto-correlation problem should affect the regression analysis provided in Section 5.2.

Table 8 reports all regression models on the expected (full) profile price. The main insight is that model coefficients are relatively stable as Resources are added in. When including bots on sale in the regression (Table 9), coefficients appear relatively stable and in line with those reported in Table 8, both in terms of trend and magnitude. An exception is $\log(GDP)$ in Model 5a, where the respective coefficient is not significant and drops in value when compared to Model 4a and Model 6a. This may suggest a correlation between $\log(GDP)$ and presence of Commerce resources for profiles on sale, that is not present or weaker for profiles at full price.

Table 6: Logistic regression for discounted profiles

$y = sale$	Model
β_0	-0.09* (0.05)
Resources	0.00*** (0.00)
Year 2020	-0.92*** (0.04)
Cookies	0.00 (0.00)
Browsers	0.14*** (0.02)
R^2	0.04
Num. obs.	11683

*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

Table 7: Autocorrelation matrix among categories of resources available for each bot.

	Crypto	Social	Services	Other	MoneyTransfer
Crypto					
Social	0.08				
Services	0.04	0.09			
Other	0.05	0.11	0.05		
MoneyTransfer	0.16	0.26	0.10	0.17	
Commerce	0.12	0.28	0.09	0.17	0.42

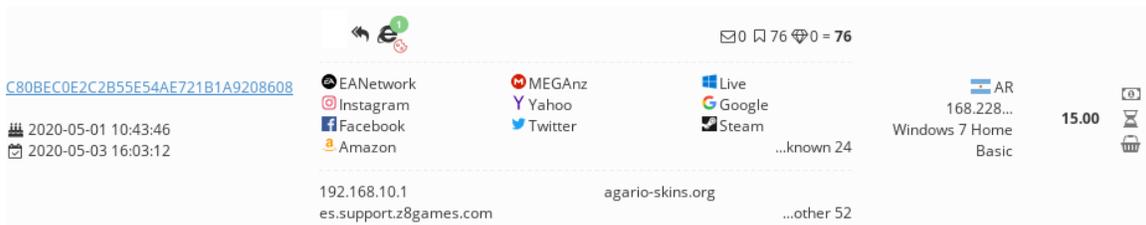


Figure 11: Overview of a listed profile on ImpaaS. ru.

Table 8: Statistical models for profiles sold at full price.

	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6	Model 7	Model 8	Model 9
β_0	8.71*** (0.08)	-12.11*** (1.21)	-11.64*** (1.13)	-2.08* (0.87)	2.24** (0.80)	-2.15** (0.80)	-5.54*** (0.82)	-5.57*** (0.81)	-3.70*** (0.63)
Real Fngrpr	1.06*** (0.14)	0.69*** (0.16)	0.80*** (0.15)	1.15*** (0.11)	1.30*** (0.10)	1.30*** (0.10)	1.30*** (0.10)	1.31*** (0.10)	1.11*** (0.07)
log(GDP)		2.29*** (0.12)	2.19*** (0.11)	0.87*** (0.09)	0.24** (0.08)	0.49*** (0.08)	0.44*** (0.08)	0.46*** (0.08)	0.42*** (0.06)
Crypto			21.74*** (0.65)	15.19*** (0.51)	14.15*** (0.46)	13.87*** (0.45)	13.72*** (0.44)	13.62*** (0.44)	10.12*** (0.34)
Money Transfer				12.30*** (0.17)	9.91*** (0.17)	9.21*** (0.17)	9.07*** (0.16)	8.86*** (0.16)	6.20*** (0.13)
Commerce					5.94*** (0.15)	5.25*** (0.15)	5.27*** (0.15)	5.06*** (0.15)	3.22*** (0.12)
Social						3.50*** (0.15)	3.52*** (0.15)	3.44*** (0.15)	1.68*** (0.12)
Services							4.08*** (0.29)	3.95*** (0.29)	2.31*** (0.22)
Other								4.22*** (0.31)	0.89*** (0.24)
Resources									0.10*** (0.00)
R^2	<0.01	0.04	0.18	0.52	0.60	0.63	0.64	0.65	0.79
Adj. R^2	<0.01	0.04	0.18	0.52	0.60	0.63	0.64	0.65	0.79
Num. obs.	7123	7123	7123	7123	7123	7123	7123	7123	7123

*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

Table 9: Statistical models for all profiles (sold at full price and on sale).

	Model 1a	Model 2a	Model 3a	Model 4a	Model 5a	Model 6a	Model 7a	Model 8a	Model 9a	Model 10a
β_0	8.74*** (0.08)	-8.76*** (0.85)	-8.29*** (0.80)	-0.17 (0.63)	3.52*** (0.58)	-0.29 (0.58)	-3.49*** (0.60)	-3.46*** (0.59)	-2.26*** (0.48)	-2.10*** (0.44)
Real Fngrpr	1.07*** (0.14)	1.22*** (0.14)	1.29*** (0.13)	1.52*** (0.10)	1.66*** (0.09)	1.69*** (0.09)	1.69*** (0.09)	1.70*** (0.09)	1.49*** (0.07)	1.03*** (0.07)
log(GDP)		1.77*** (0.09)	1.69*** (0.08)	0.58*** (0.06)	0.03 (0.06)	0.25*** (0.06)	0.21*** (0.06)	0.21*** (0.06)	0.24*** (0.05)	0.36*** (0.04)
Crypto			19.23*** (0.49)	13.76*** (0.39)	12.72*** (0.36)	12.48*** (0.35)	12.34*** (0.34)	12.25*** (0.34)	9.08*** (0.28)	8.91*** (0.25)
Money Transfer				10.95*** (0.13)	8.82*** (0.12)	8.25*** (0.12)	8.12*** (0.12)	7.94*** (0.12)	5.58*** (0.10)	5.37*** (0.09)
Commerce					5.21*** (0.11)	4.60*** (0.11)	4.60*** (0.11)	4.42*** (0.11)	2.71*** (0.09)	2.74*** (0.08)
Social						3.02*** (0.11)	3.08*** (0.11)	3.01*** (0.11)	1.42*** (0.09)	1.50*** (0.08)
Services							3.83*** (0.21)	3.72*** (0.21)	2.13*** (0.17)	2.17*** (0.16)
Other								3.54*** (0.22)	0.59** (0.19)	0.69*** (0.17)
Resources									0.10*** (0.00)	0.09*** (0.00)
Sale										-3.40*** (0.07)
R ²	<0.01	0.04	0.15	0.49	0.57	0.59	0.60	0.61	0.74	0.79
Adj. R ²	<0.01	0.04	0.15	0.49	0.57	0.59	0.60	0.61	0.74	0.79
Num. obs.	11683	11683	11683	11683	11683	11683	11683	11683	11683	11683

*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$