Exploring the Genesis Supply Chain for Fun and Profit

ke-la.com/exploring-the-genesis-supply-chain-for-fun-and-profit/

February 21, 2020

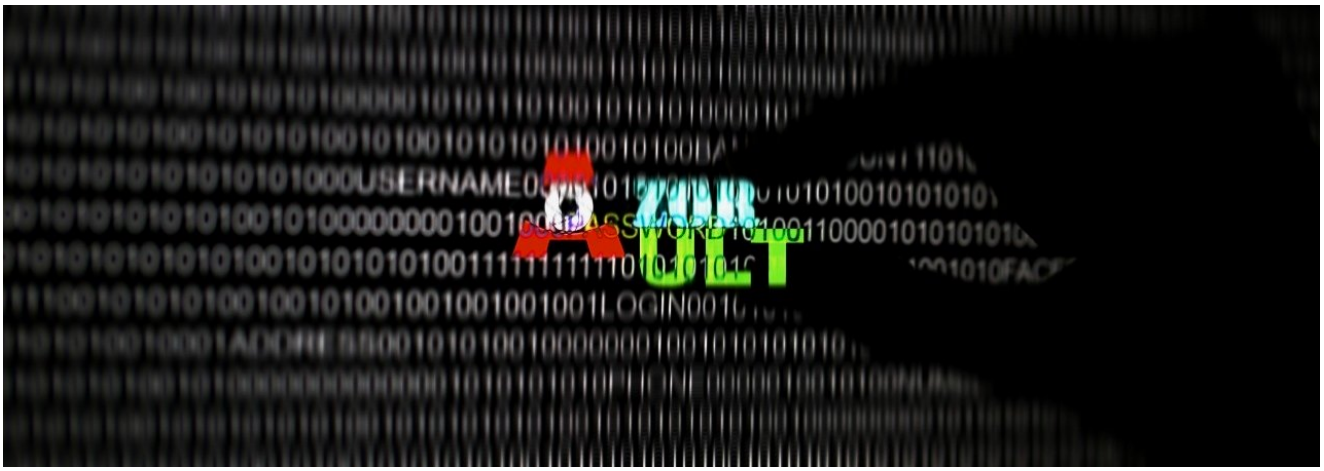

Bottom Line Up Front

This is the first post in a series of posts reviewing the supply chain of the Genesis Store market – a likely-Russian threat actor operating a successful, borderline innovative, *pay-per-bot* store since 2018. The following post features a quick-and-easy methodology breaking down over 335,000 unique Genesis infections into four malware groups, allowing us to attribute **over 300,000 AZORult infections to the Genesis actors currently involved in campaigns resulting is tens of thousands of new AZORult infections per month**. Furthermore, it seems Genesis isn't necessarily leading these campaigns, but rather working with various Malware-as-a-Service (MaaS) providers and cybercrime services.

This discovery, linking Genesis with widely known commodity malware, **highlights the ongoing threat to organizations and the proliferation of illegal data obtained from infections**. It also sheds light on the supply chain relationships between actors operating within the cybercrime financial ecosystem (read: Dark Net); we'll explore this theme, including specific actors and trends, in the next posts.

Preface

Much has been written about the Genesis Store: the threat it poses to organizations, the native fingerprinting abilities embedded in the tools it provides, or (shameless plug alert) the shift in cybercrime business models it represents.

On a day-to-day operations basis, **KELA automatically monitors new listings on the market** – allowing our clients to monitor relevant infections. While these real-time alerts allow threat intelligence or incident response teams to remediate relevant threats, many of our clients were interested in a wider, more contextual understanding of Genesis. For example, one report indicated that Genesis *"introduced a new breed of stealers specifically designed to collect digital fingerprints and artifacts;"* another referred to the Genesis *"botnet."* That raises the question, **what do Genesis actors actually do**? Develop specialized stealers, manage a MaaS botnet, or sell credentials obtained by other actors?

Seeking to demystify the threat, we focused on three major areas:

1. Technical abilities – which trojan is being used: commodity malware or a tailored tool?

2. Operational independence – are the Genesis actors leading independent campaigns, determining targets and infection methods, or reselling data obtained by other parties?

3. Scale – the spread, infection rate and profits pocketed by the actors

Our research aims to evaluate the actionable threat level: if ACME Corp found an infected endpoint offered for sale on the market, what should be the next steps of the incident response team? What does a Genesis infection say about further threat hunting, threat containment and lateral movement, or impending payloads? In the age of cybercrime inter-group relations and joint campaigns, understanding how Genesis actors position themselves in the cybercrime financial ecosystem as an emerging threat group can be extremely helpful to defenders.

Luckily, it's our business to have access to large cybercrime-related data; as such, **KELA's systems have been scraping Genesis ever since it became a prominent market,** caching every infected machine's data and metadata, resulting in over 335,000 unique devices at the time of writing. However, unlike a day-to-day malware analysis assignment, **we had no specific infection to investigate**. Since KELA is an external service provider, we learn about Genesis infections affecting our clients from listings on the market and any network or endpoint indicators, leaving us with no access to actual samples.

So, the first question to answer was:How do you analyze a campaign without traditional indicators of compromise?

Let Me GUID You

Let's begin by introducing the components of a typical Genesis listing: the bot metadata, providing context and data on the infected machine; and the data itself – showing the actual stolen credentials and associated details.

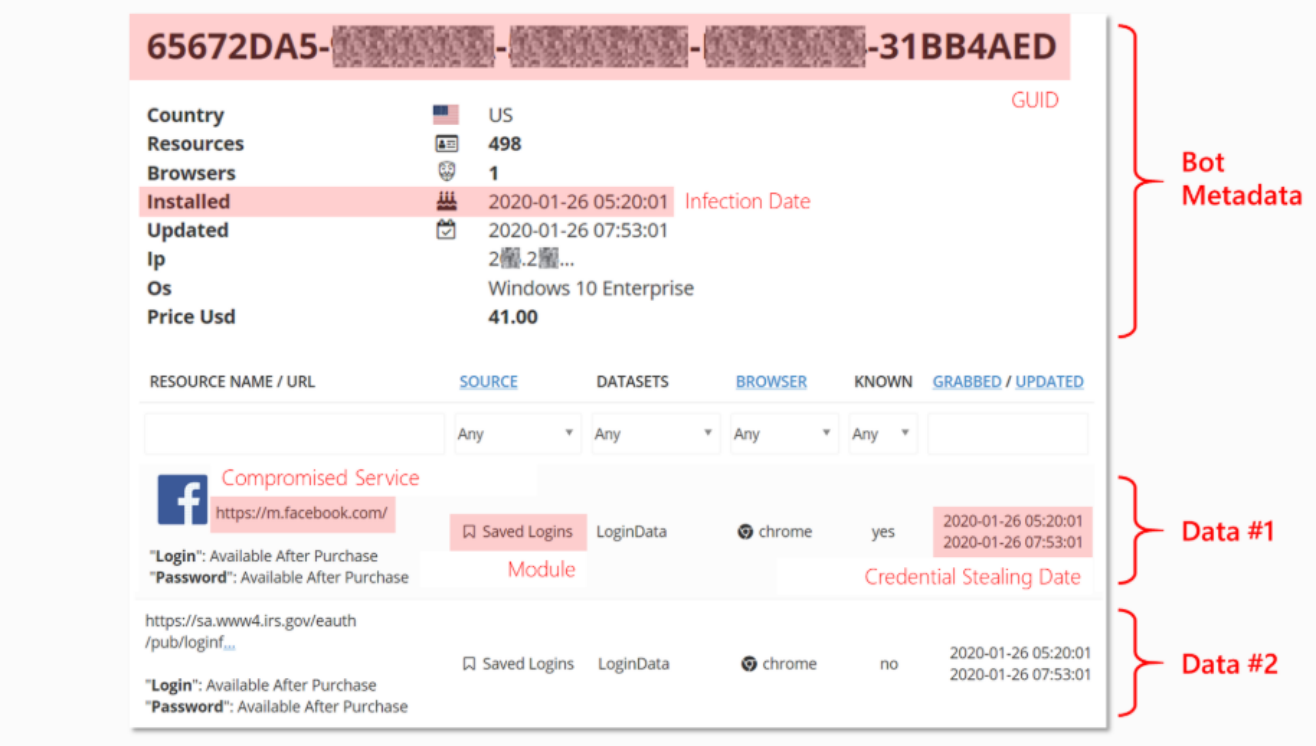

Figure 1: Typical bot listing on the Genesis market

Notably, Genesis obscures a significant property, which other actors selling infostealer data provide: the infection type. While other markets and sellers clearly indicate which Trojan was used to source the stolen credentials, Genesis prefers white-labeling and rebranding data as their own.

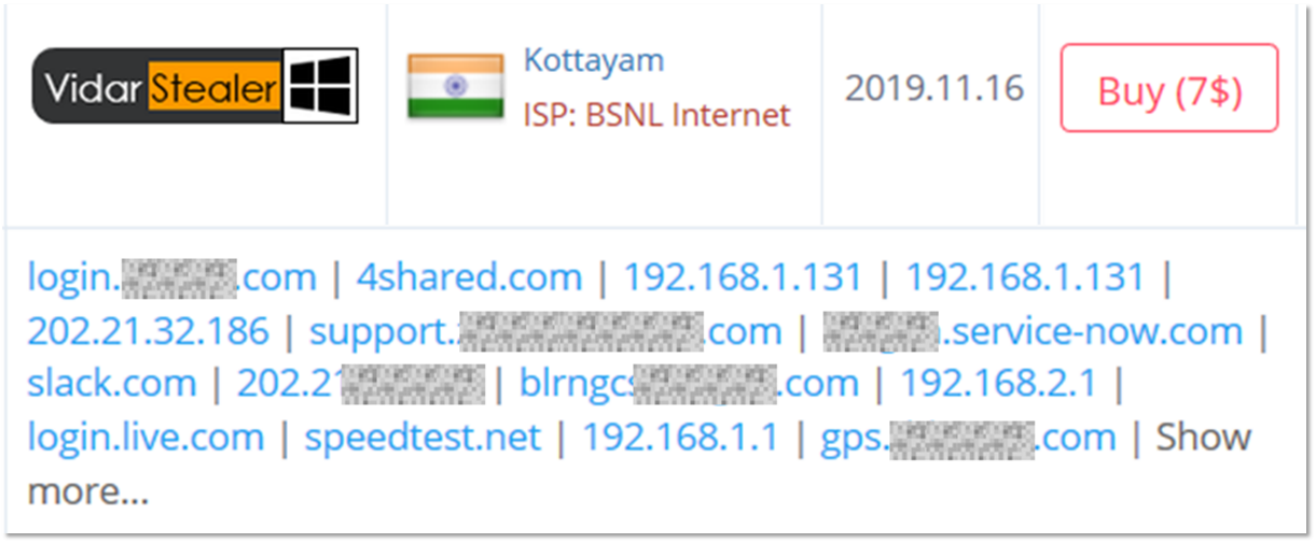

Figure 2: A bot for sale on a different market, stating it was obtained using Vidar stealer

---

With our **main area of interest being the attributed stealer** – the one data point Genesis actively hides – our methodology focused on trying to extrapolate the available data points. The *GUID, module* and *infection date* proved useful in answering that question.

Our preliminary assumption is that the *bot name* – the title of each Genesis bot – **is a globally unique identifier**, which we will refer to as the bot GUID. To put that into context, note the structure of the bot GUID in Figure 1: it's composed of five alphanumerical octets separated by hyphens. However, not all Genesis bots follow this distinct pattern; scanning the website, we've identified several GUID structures. So, **if each GUID structure is associated with a particular malware, could we use the patterns to classify infection types**?

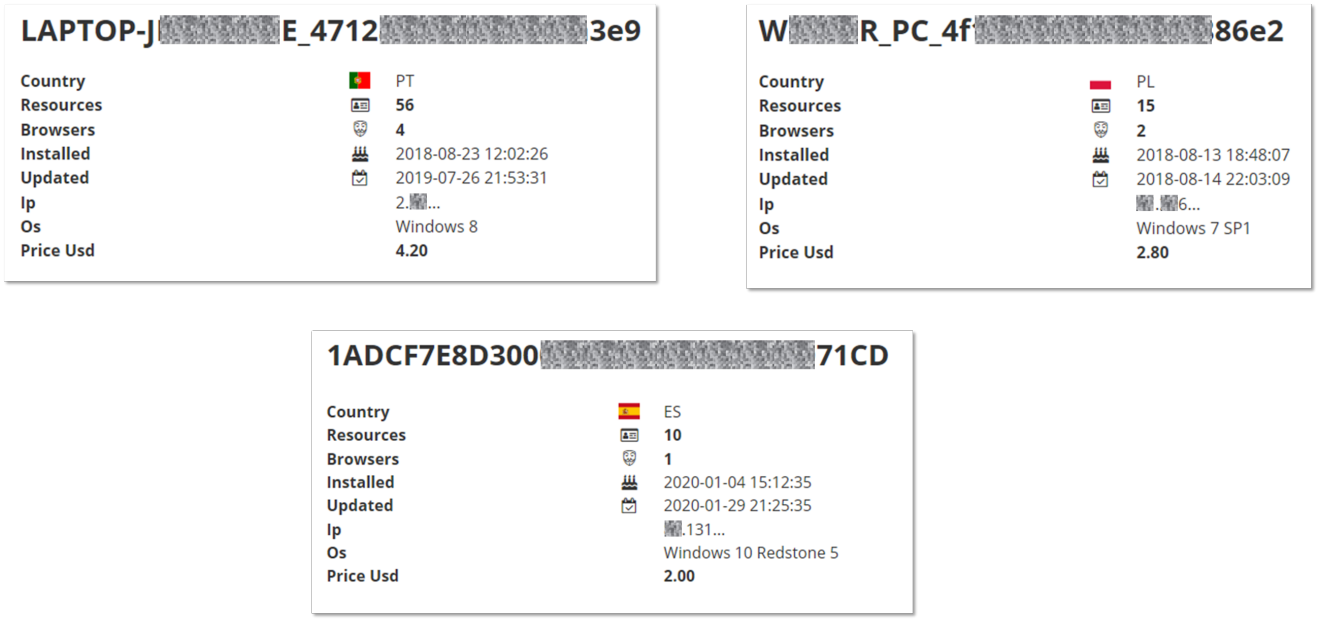

Figure 3: Genesis bots with different GUID patterns

To answer that, we attempted to **group all >335k Genesis bots scraped by KELA into distinct *GUID classes***, based on the structure and syntax observed. We then assigned every GUID group to a malware, effectively mapping the Genesis malware supply chain. Our methodology involved three steps:

1. Mapping possible delimiters and splitting each GUID into **tokens**, accordingly;

2. Mining seemingly **meaningful strings** from all tokens into a dictionary; and

3. Classifying each GUID based on the **length of the token** in its GUID and whether an important string appears in any of the tokens.

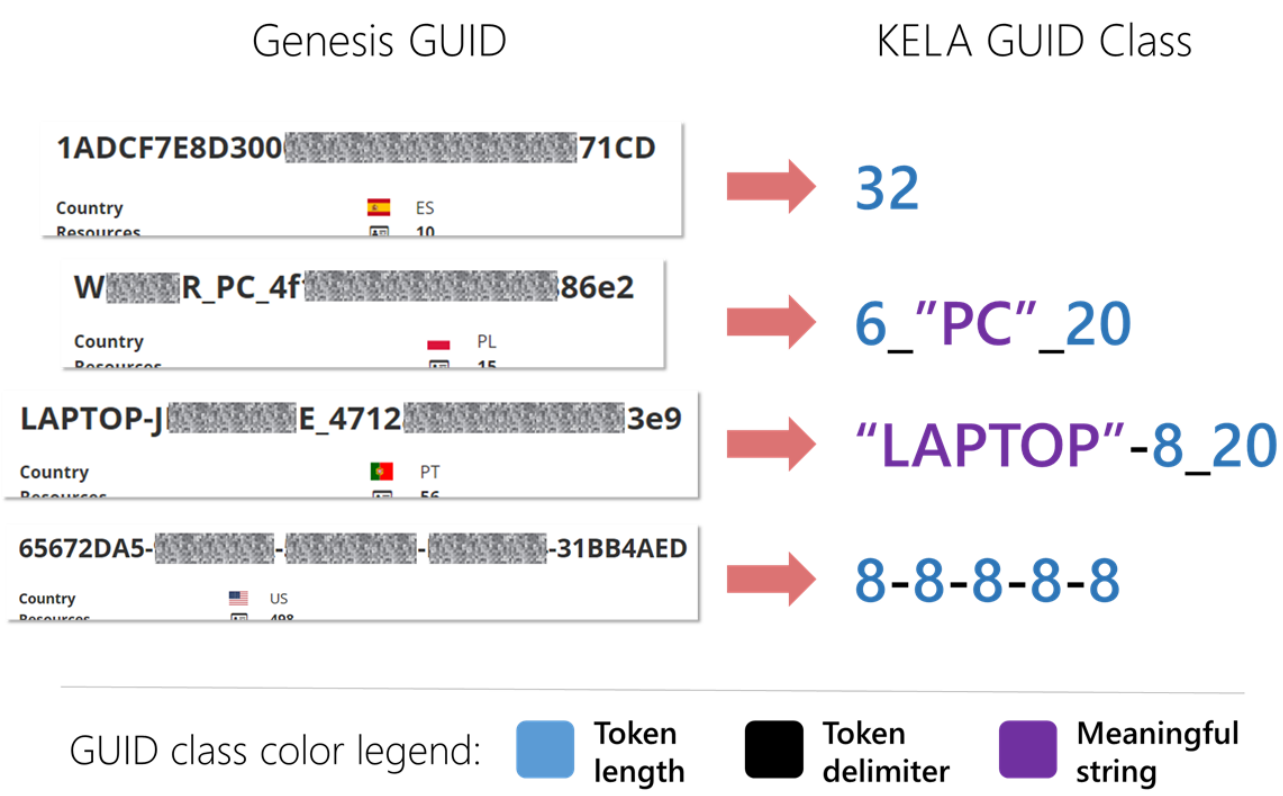

Figure 4: KELA's classification of different GUID patterns in Genesis

---

Incorporating this methodology, we grouped all Genesis bots into 340 distinct GUID classes. That might seem like a major achievement, allowing us to map 340 different GUIDs to known (or unknown!) information stealers. However, as can be seen in Figure 5, many of the GUIDs share another syntactic feature: they end in "_20," meaning that regardless of how the raw GUID starts, it ends in 20 alphanumeric characters following an underscore.

Class	# of Tokens	Relative Part
8-8-8-8-8	5	90.3856%
32	1	5.6381%
“DESKTOP”-7_20	3	1.3892%
8-4-4-4-12	5	0.2568%
“LAPTOP”-8_20	3	0.2435%
4-"PC"_20	3	0.1956%
5-"PC"_20	3	0.1635%
6-"PC"_20	3	0.1573%
4_20	2	0.1304%
6_20	2	0.1276%

Figure 5: Initial GUID classes: the grey ones are distinct and well-formed, while the pink ones are less distinct, and were later grouped based on secondary properties

The Four GUIDers of the Apocalypse

This insight, along with the meaningful strings appearing in some GUIDs, quickly led us to realize that these GUIDs are essentially a Windows machine name concatenated with a seemingly unique 20-character string. We refer to this class as *{MACHINENAME}_20*. This tweak to the classification methodology allowed us to group over 330 different GUIDs, resulting in **four key structures across the entire Genesis dataset.**

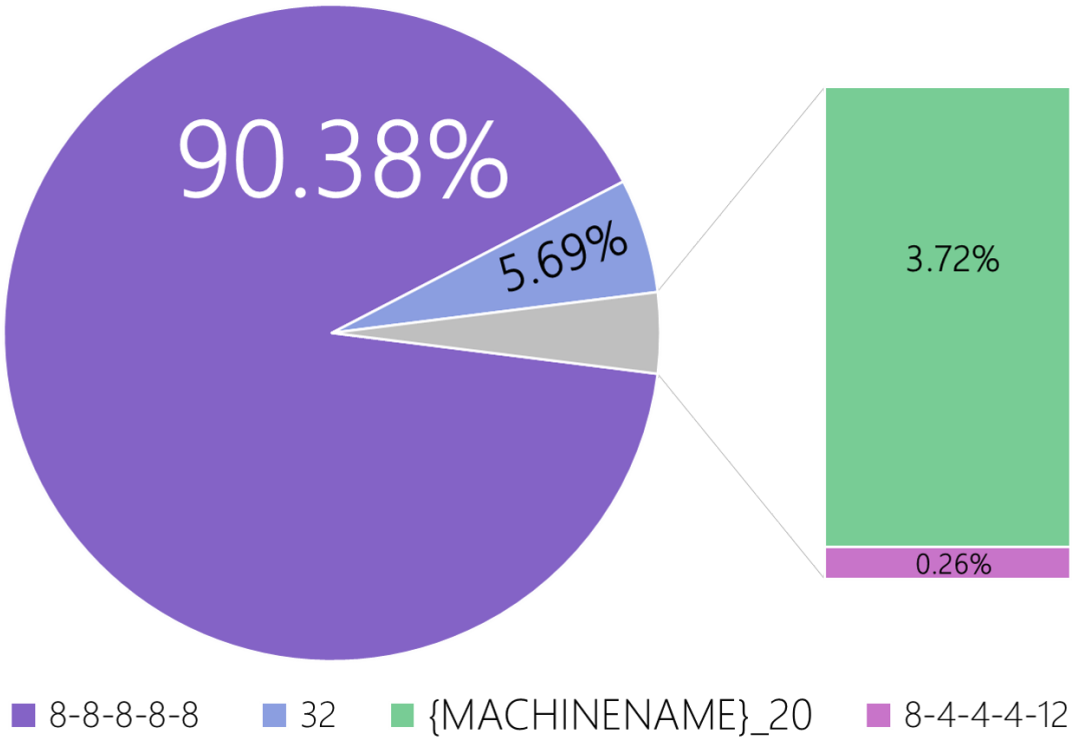

Figure 6: Breakdown of bot GUID structures

With the four malware families classified and the number of infections accounted for, we could decipher other meta-properties and make sense of how Genesis operates. Reviewing infection classes over time suggests **Genesis actors have been experimenting with different Trojans**: as seen in Figure 7, during its first year, Genesis only sold infections of *{MACHINENAME}_20* class, and switched to *8-8-8-8-8* in late 2018, boosting the number of infected machines obtained by the actors.

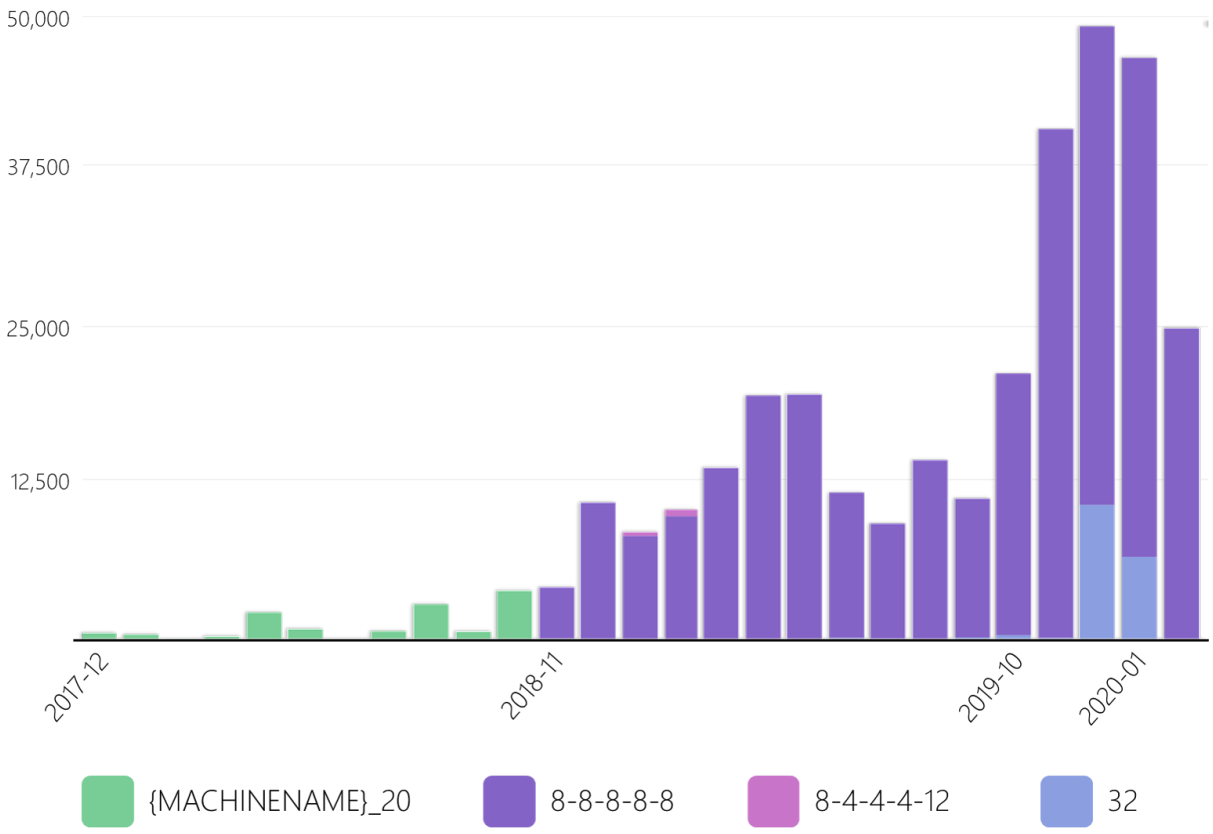

Figure 7: Infection types (by GUID class) over time; dates indicate when the machine was infected

---

Another interesting date was late 2019 when the *32* GUID class appeared on the market; Genesis actors might have considered introducing a new malware type into the market, but eventually decided to revert to their leading product – the 8-8-8-8-8 class.

Armed with these insights, **we set to analyze the *8-8-8-8-8* class** based on its ongoing ubiquity. In our next post, we'll dive into the *32* class, providing insights into its short rise and fall on the market, and explore the historical themes of the two remaining categories.

88888, the Number of the GUID

Seeing this pattern, **our first suspect was AZORult infostealer**, a malware that uses a similar GUID. However, it was circumstantial at best; we needed some hard evidence linking the two.

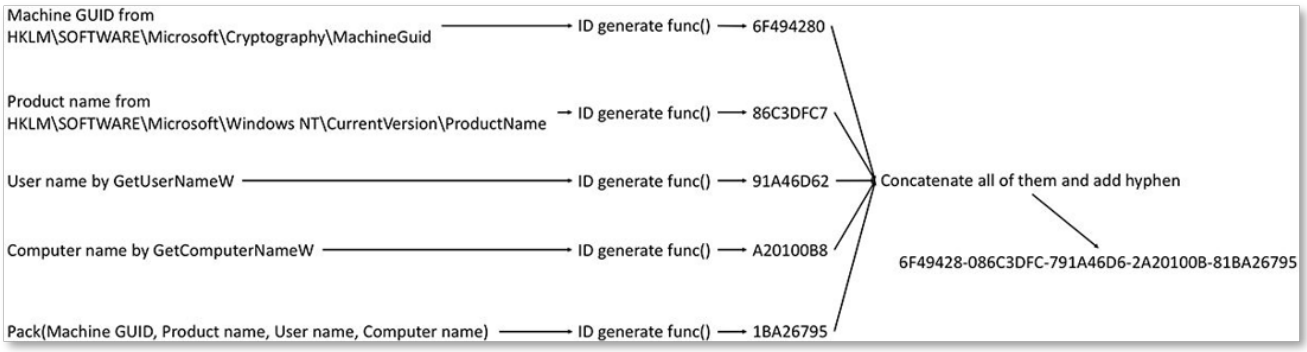

Figure 8: AZORult's GUID featuring the 8-8-8-8-8 structure (analyzed by Cylance)

Luckily, the cybercrime financial ecosystem is abundant with actors sharing samples of their "logs" – collections obtained via various malware types. We soon came across a prominent MaaS provider operating in a top-tier Russian cybercrime forum. Retrieving some of the data shared by the actor, **we were able to find official AZORult logs** – stolen credentials advertised as having been obtained via AZORult – **that share the same GUIDs as Genesis' *8-8-8-8-8* category**. That, along with corresponding metadata, confirmed the *8-8-8-8-8* class, which comprises over 90% of Genesis infections, is based on AZORult.

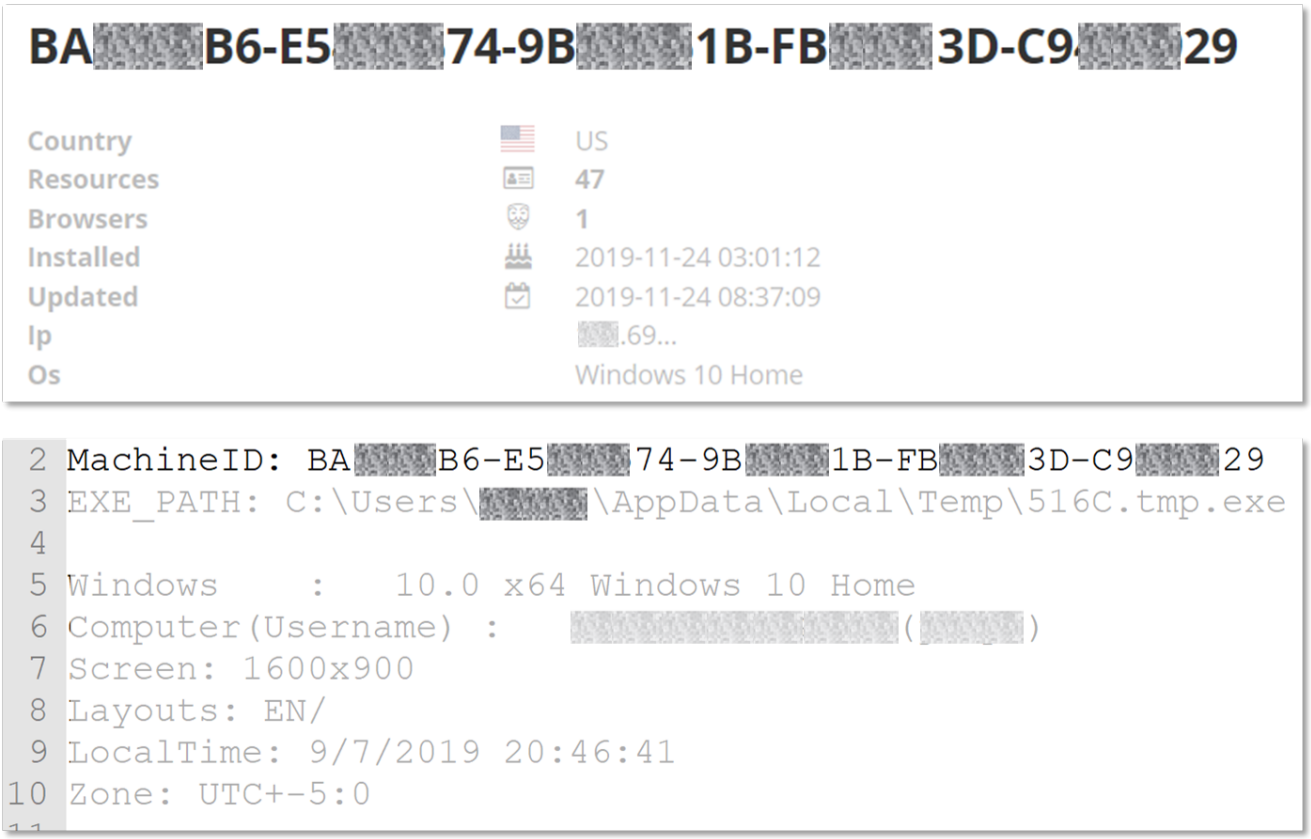

Figure 9: Genesis listing (top) and an AZORult log obtained from a prominent MaaS provider (bottom) sharing the same GUID and metadata, confirming the "8-8-8-8-8" Genesis GUIDs are indeed AZORult

Operating since 2016, AZORult is a commodity malware widely used by multiple threat actors in numerous campaigns. While its author stopped maintaining the project in late 2018, AZORult is still widespread and used in active campaigns. AZORult's source code was readily available to numerous actors who modified the original, producing unofficial variants for small-scale, independent campaigns. In the case of Genesis, this could mean actors either operate their own version of the stealer, resell infection from a MaaS provider, or both.

Going back to the AZORult log samples mentioned in Figure 9, we analyzed fifteen months' worth of records offered by the MaaS provider and found correlating GUIDs and **over 10,000 such bots**. These findings might indicate a **supply chain link between Genesis and known cybercriminals**, where actors monetize their campaigns by selling data to Genesis.

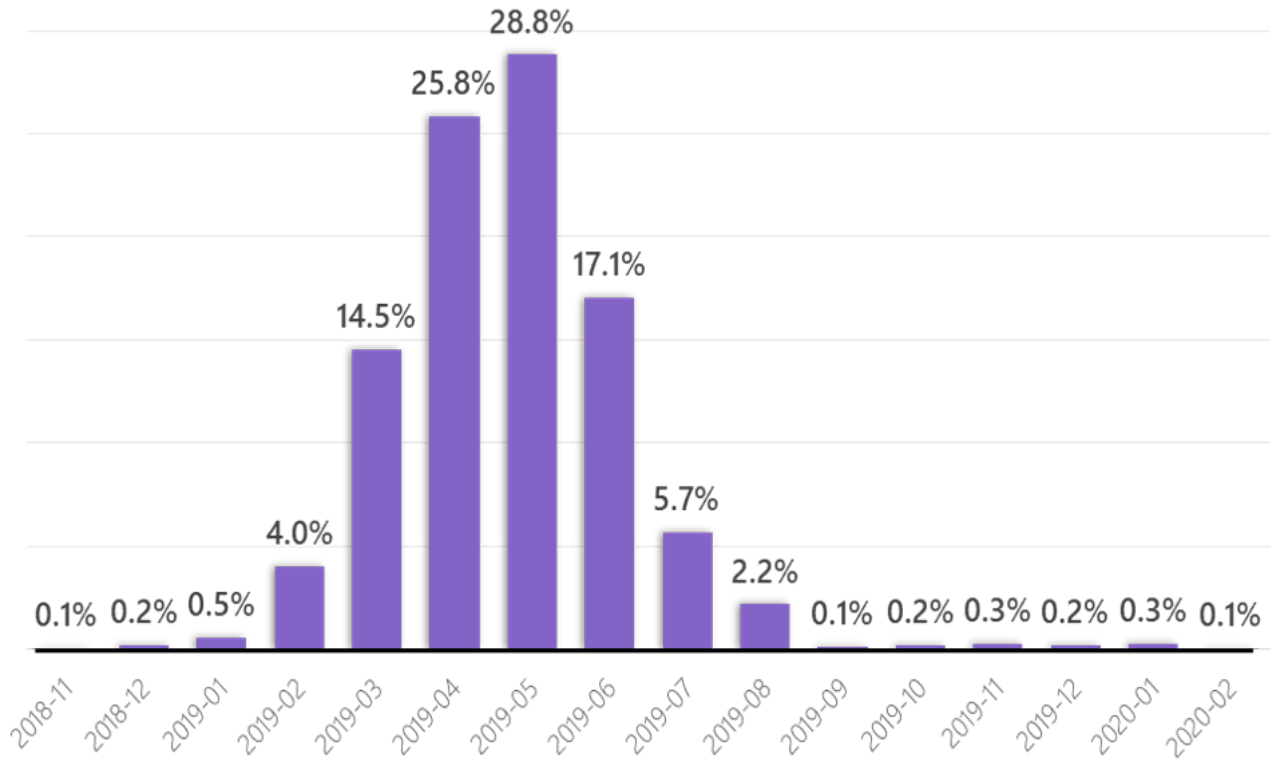

Figure 10: Possible cooperation between Genesis and a MaaS provider offering access to AZORult-infected machines

Not only does Genesis use commodity malware for over 90% of its bots – the actors do so as part of a joint business venture with MaaS providers. This offers a glimpse into an interesting aspect of the cybercrime sector: interactions and cooperation between actors. In the case of Genesis, **old forum posts indicate they're not entirely independent and are**

interested in gaining access to compromised machines. As can be seen in Figures 7 and 10, Genesis obviously found a supplier – most likely several – of infections they can utilize for credential-stealing purposes.

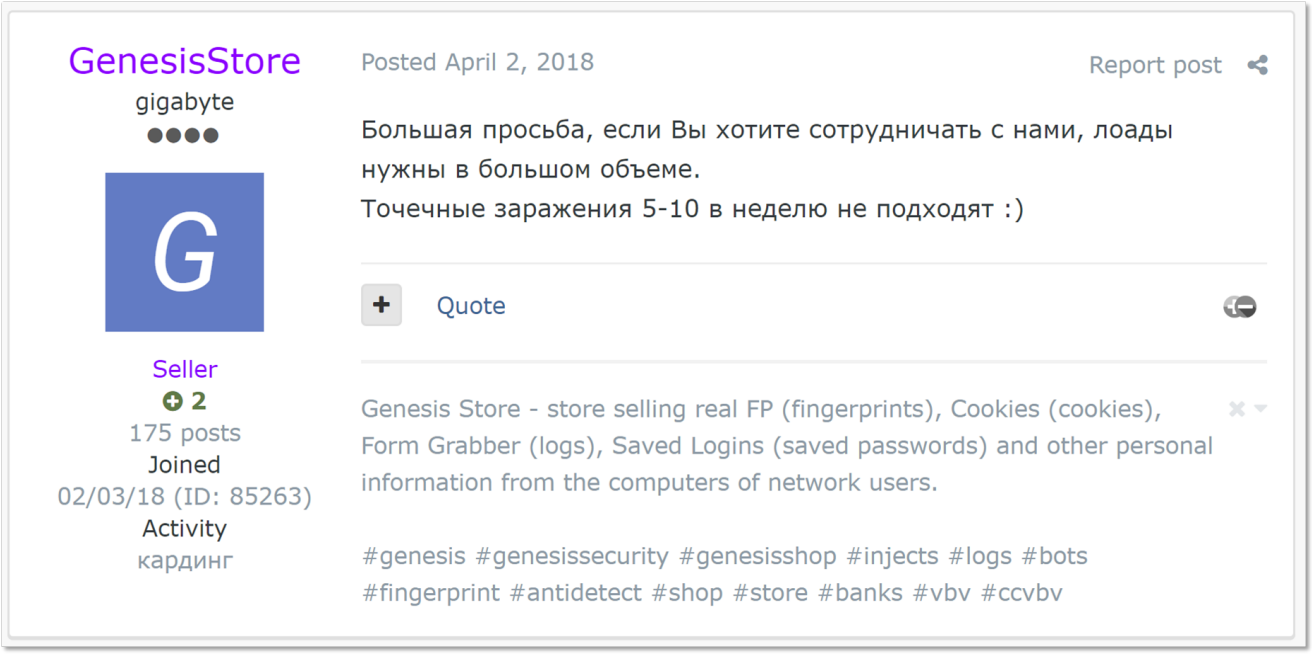

Figure 11: A post from the early days of the Genesis Store: "Asking you kindly, if you want to work with us, we need loads in big volumes. Topical infections of 5-10 a week will not do"

The next parts of this research will dive deeper into the supply chain relations Genesis has with at least one prominent MaaS provider, shedding light on cybercriminal business models.

The GUID, the Bad, and the Obvious

In the first part of our research, we focused on developing a methodology for continuous monitoring of the Genesis Store – specifically supply-related trends. Analyzing the elements of one of today's most prolific cybercrime outlets, **we were able to link over 300,000 AZORult infections to Genesis** – or at least to their providers. Behind these numbers are diverse victims: from SMBs to enterprises, from the private sector up to government officials.

While the fact that a top-tier market uses one of the most popular stealers is hardly surprising, our main point is this: **Genesis actors are not going anywhere anytime soon, so we might as well keep tabs on them.** Gaining a better understanding of the threat, as well as demystifying it, is crucial in establishing KPIs for further monitoring. Only now – when we have clear definitions of the necessary data points to extrapolate from the operations of Genesis actors – can we establish baseline metrics and measure when they're being disrupted. Adding new malware strains, utilizing new and different credential-stealing modules, or shifting infection volumes can serve as threat indicators to improve awareness.

This is also our motivation for the upcoming parts of this research, where we will try to:

- Identify the three remaining GUID classes, linking them to known malware;
- Describe the different stages of Genesis' evolution, from inception to a top-tier provider of the cybercriminal underground;
- Deep dive into specific MaaS actors involved, both directly and indirectly, in the Genesis supply chain;
- Explore themes and trends in the MaaS industry and the broader cybercrime financial ecosystem, and their effects on consumers of threat intelligence;
- Refer to biases, traps and gaps in our methodology; and
- Come up with as many GUID-themed puns as possible.