NETACEA

Technical Showcase: Netacea's Approach to Machine Learning in Advanced Bot Management

INTRODUCTION

The bot threat landscape is getting more complex all the time. The prevalence of bot traffic on the web is growing in line with society's enthusiasm for digital services, which has only increased due to the Covid-19 pandemic.

Netacea research estimates that, on average, 53% of all web traffic is automated rather than directly human. Not all of these bots are bad, but malicious bots account for 33% of all web traffic and can cause untold damage to websites and their users.

In previous years, it has been relatively simple to uncover malicious bots by looking at their origin or device profiles. However, the creators of these bots are going to great lengths to avoid detection.

To make matters even more complicated, there is a huge variety of business logic attacks being carried out by bots. Depending on their goals, attacks vary in:

/ The volume and velocity of traffic

/ The origin and distribution of requests

/ The specificity of the paths they target

NetBLADE: Netacea Business Logic Attack Definition framework

As the bot attack landscape grows in complexity, Netacea saw a need for more consistency around how bot threats are classified. We have collaborated with cybersecurity influencers across the industry to create NetBLADE, an open-source framework to capture the stages, tactics and techniques used in business logic attacks of all kinds.

With so much disparity between threats, a multi-layered strategy is needed to detect and mitigate all bot attack types across the landscape.

Find out more about NetBLADE: Netacea Business Logic Attack Definition

OUR APPROACH: BOT EXPERTS AND DATA

The more sophisticated bots are now able to disguise themselves as real users through spoofed device fingerprints, geolocation, IP addresses and other signals. At Netacea we root out bot traffic by analyzing user behavior on the server side.

We do this by applying valuable domain knowledge and artificial intelligence (AI) to web log data. Regardless of their objective, bots have a range of common behaviors that can be detected through machine learning and analytics to distinguish the intent of every request.

Identifying common threat behaviors

To ensure our approach encapsulates every kind of business logic attack, we first call on our expertise in the field to identify the different and often opposing characteristics of bot attacks. These traits aid the bots in carrying out their attacks, but the behaviors are what make them susceptible to detection.

Our aim is to use machine learning to capture as many bot behaviors as possible through a variety of tools and techniques. Here are some of the more well-known bot behaviors:

Volume and velocity

Traditional bot defenses look for high volumes of traffic making similar requests in a short period of time. This kind of high-volume approach is necessary for credential stuffing or account takeover attacks, as even a small success rate can yield big profits.

However, other bots seek to avoid discovery by adopting a "low and slow" approach, hitting their target at volumes low enough to avoid detection thresholds whilst still causing significant damage over time. One example is scraping or data harvesting, which could be running continually but not at volumes that would indicate abnormal activity.

Origin

The origin of a bot attack is often dependent on the resources available to the bad actor orchestrating it. Many requests coming from a single source, especially from a known bad data center or a geolocation unlikely to be genuine human traffic, is an obvious bot attack, but one common enough to require close attention.

More sophisticated attacks mask their true origin by distributing requests across IP addresses. These attacks are especially difficult to track because they often originate from home ISPs in unsuspicious geolocations. Most commonly these attacks are orchestrated via botnets to use infected home devices.

Target

The target of a bot attack is usually very specific. This could mean only one path is hit repeatedly, making detection simpler. For example, card cracking attacks will only be interested in the checkout page.

To disguise themselves, some bots will add multiple or varied request paths into their journeys. More complex objectives may also require multiple targets to be hit, especially if the bot is carrying out several stages of a larger attack.

The intangible behavior profile – Humans

While website owners design user experiences and try to predict how website visitors will act, in truth the behavior of real humans is hard to anticipate. Yet, it is important to distinguish what is human so that legitimate users are not blocked erroneously.

DATA SCIENCE CHALLENGES

With such variation in behaviors, training a simple classifier for each is not a viable approach.

The data we glean from web logs is often hard to label appropriately. We may not have the right label available for a set of data, and the labels we do have could be unbalanced or "fuzzy" – potentially too broad to offer value.

Data can also suffer from "concept drift", where the website being monitored can change both gradually over time or dramatically overnight. The behaviors of visitors can also change from usual expectations, for example during events like Black Friday in retail.

The execution of our bot mitigation strategy depends on the client's appetite for risk, both in speed of blocking and in setting thresholds between what is deemed malicious, benign, useful, or human behavior.

Another challenge in using machine learning to guide bot mitigation is managing "unknown unknowns". There are always new domain architectures appearing, and new attack vectors to contend with, making bot mitigation a constantly moving target.

APPLYING MACHINE LEARNING TO ADVANCED BOT MANAGEMENT

At Netacea we look to utilize a variety of tools and techniques to capture different behaviors; it is useful to explore the pros and cons of some different approaches before diving deeper and looking at how we can apply this thought process to a specific problem.

DIFFERING APPROACHES TO DETECTING BOT THREATS

Real-time vs. batch

At a very high level, real-time processes are challenging yet can be extremely effective at stopping high-risk activity such as credential stuffing or carding attacks.

On the other hand, batch processes can be more precise, offer more context and require less sacrifices for speed. Batch processes can be used to provide additional confidence or to stop behavior with a longer time to value, like scraping attacks.

Unsupervised vs. supervised

Likewise, we look to use a blend of supervised (label trained) and unsupervised approaches.

Supervised approaches are easy to evaluate and scientifically rigorous, however they require robust labels and a well understood data set to be effective.

Realistically we must rely on unsupervised methods to track dynamic behavior across domains, especially for real-time attacks. Although difficult to evaluate or monitor, unsupervised methods are more adaptive and can make use of more data, especially as this does not need to be labelled data.

General vs. specific

Finally, we look to find a balance between general approaches, which allow us to identify bots across a range of situations, and specific methods.

Our techniques range from general scraping models and intelligence feeds, through to specific models to help in very precise business logic attacks.

General models are adaptive and can make an immediate impact on new domains. They are also easier to maintain. However, they are not always ideal for specific business logic. This is where we need specific models, though it isn't feasible to have a model for every single case, especially as there are many different stages and behaviors within each attack.

Considering all these approaches gives us eight distinct overall methods, each with their own advantages and disadvantages. In reality, not every attack fits into just one category and there can be crossover, but it is useful to lay this high-level view over the NetBLADE framework to identify where different approaches can work together.

Supervised		
	General	Specific
Realtime	1	2
Batch	3	4
Unsupervised		
	General	Specific
Realtime	5	6
Batch	7	8

Fig. 1 shows the differing approaches to machine learning in advanced bot detection.

EXAMPLE IN ACTION: SCALPER BOTS

To demonstrate how different approaches can be applied across a "kill chain" as identified by the NetBLADE framework, we can turn to a well-known example: scalper bots.

Scalper bots use automation to hit retail and booking sites and buy up high-demand stock before human users have the chance to purchase. These low-supply, high-demand items are then put up for sale on secondary markets at an increased price.

Scalper bots have grown in prominence over the last year as the pandemic caused physical stores to close, making many sought-after goods and services exclusively available online.

This has created opportunity for scalpers to make huge profits, preying on the desperation of consumers wanting to buy items such as PS5s, gym equipment or even grocery shopping delivery slots and Covid vaccinations.

It's not just professional scalper groups who have contributed to this. Amateurs and individuals have also got involved in scalping, either by renting or buying bots from professionals or by writing their own scripts, to get their hands on wanted items or to make some quick money of their own. Such activity is closely tied to "hot drops" of popular items and is not technically illegal.

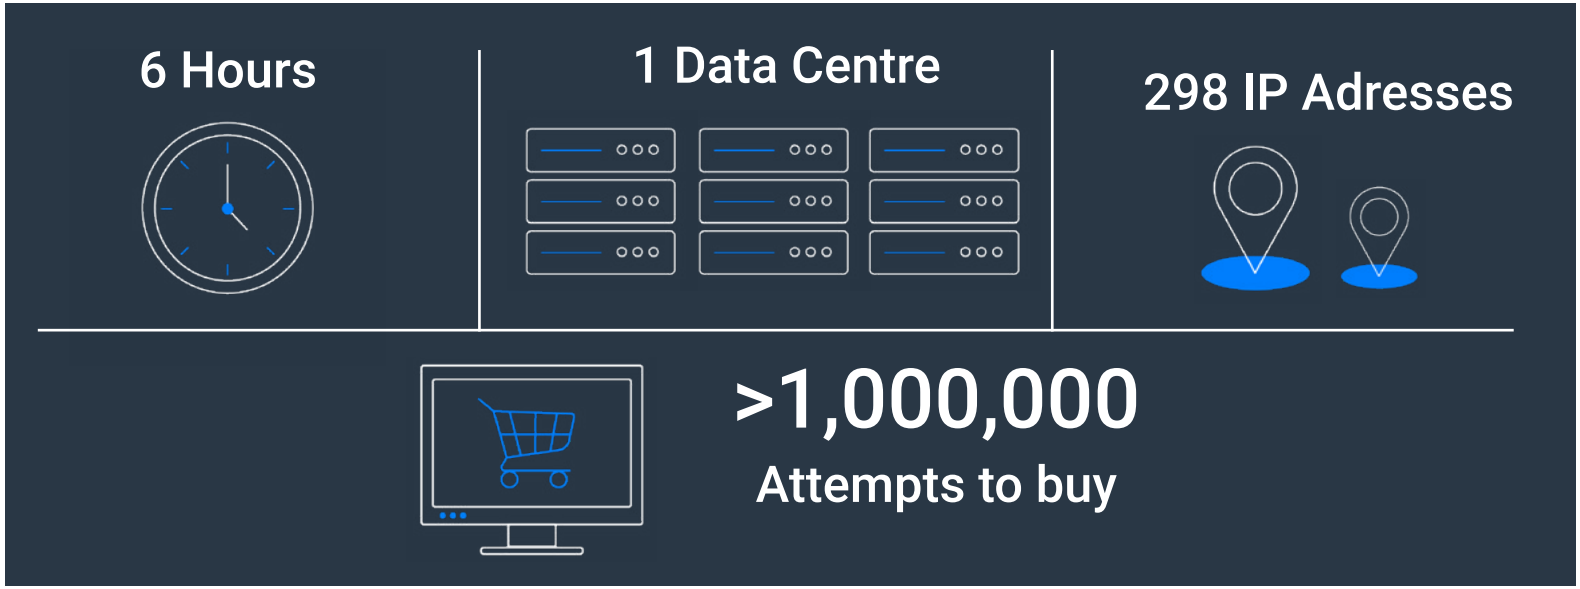

Fig. 2 shows scalper bot attack data where over a million attempts to buy were made in a six-hour period

INSIDE THE SCALPER BOT KILL CHAIN

Thinking about the journey of a scalper bot, we can see various distinct stages to the process. This chain might span several weeks prior to and after the "drop". Using the NetBLADE framework, we can identify key parts of this chain where we can most effectively stop the bad actors, and feed any learnings into a reputational database to use against other attacks.

1 Resource development

Every bot attack requires infrastructure, and we are often able to detect when such infrastructure is being tested ahead of a bigger attack. We can use a combination of real-time and batch processes to spot distinct tests across domains, often using historical reputational knowledge to highlight bad actors early on.

2 Attack preparation

While we can often block resources early, these are often cheap enough for bot operators to re-tool and move into attack preparation. This involves activities such as account creation, account takeover, and account aggregation, where scalpers will build up many accounts on target websites with which to make their purchases. This is where Netacea's real-time behavioral tools come into play.

These attacks are typically high volume to maximize their success, but this makes them identifiable to real-time behavioral analysis.

These behaviors are also fairly generic, although we may need to use specific machine learning approaches to root out certain types of business logic at this stage.

3 Reconnaissance

In this phase, the attackers pick their target, whether this is a loose cluster of victims or a specific target, and monitor them for changes to detect the specific second that items go on sale.

This often sits alongside other attacks and involves long-term monitoring. As the time to value is longer, we find most value in a batch approach to this data.

There are lots of examples of this kind of behavior, making supervised machine learning extremely effective in stopping bad actors in the reconnaissance stage.

4 Attack execution

Even after weakening the attack in the previous stages, it's likely there will be bots present during the "drop", trying to obtain the product as quickly as possible. Here we rely on a combination of real-time and batch approaches as well as utilizing all the intel further down the kill chain to protect the site.

CONCLUSION

Bot attacks are sophisticated, multi-stage processes. Each stage of the attack can be undertaken through different behaviors.

A combination of different data-driven approaches in conjunction with domain expertise is needed to stop bots from achieving their ultimate goals, as well as the ability to adapt and react to emerging threats across the bot landscape.

Machine learning is highly effective for this purpose, but only when applied in a considered way and blending different approaches where necessary.

TAKE THE RIGHT APPROACH TO YOUR BOT PROBLEM

Netacea delivers the next generation of bot management technology. Our revolutionary server-side approach and innovative Intent Analytics™ engine, powered by machine learning, allows us to analyze every request and detect a wide variety of bot behavior.

To find out more about Netacea's unique approach to stopping sophisticated bot threats, visit www.netacea.com/why-netacea or talk to our team today at hello@netacea.com.