EBOOK

The Anatomy of Account Takeover Attacks

Contents

What is Account Takeover?

Ac-count take-o-ver
noun

A type of identity theft where a bad actor gains unauthorized access to an account belonging to someone else. Also known as brute force login, dictionary attack, credential stuffing, or credential cracking.

OWASP Definitions

To understand the analysis, we define a few terms for consistency. OWASP, in its Automated Threat Handbook, defines two automated threats that use credentials. Note their distinction:

- **Credential stuffing:** Mass login attempts used to verify the validity of stolen username/password pairs.
- **Credential cracking:** Identifying valid login credentials by trying different values for usernames and/or passwords.

THE PURPOSES OF ACCOUNT TAKEOVER:

- Test sets of credentials for validity, then sell validated pairs on the dark web.
- Gain access to an account and information available therein (e.g., stored credit card data, personal information, and PII) in order to sell, distribute, or use that data.
- Potentially use an account for personal gain (such as transfering money, purchasing goods, spreading an agenda, or other website functions).

Methodology

The 2018 Bad Bot Report released by Distil Networks (acquired by Imperva), reveals that bad bots are found on websites with login pages. Since login pages are one of the most abused pages on a website, we went on to study the anatomy of account takeover attacks in greater depth in 2018. In this subsequent study, we evaluated data from 600 domains that include login pages, then analyzed a subset of 100 domains that have the largest bad bot traffic datasets.

Every Login Page Has a Bad Bot Story

In this account takeover study, Imperva now sees bad bot traffic on 100 percent of all monitored login pages. This indicates that every website is being hit by account takeover attempts.

What Do Account Takeover Attacks Look Like?

Referring to the OWASP definitions, account takeover attack attributes are split into two groups. In our study, the split was about even; half were targeted volumetric credential stuffing attacks and half were the low-and-slow credential cracking and credential stuffing attacks.

Account Takeover Types

The Difference Between Volumetric and Low-and-Slow Attacks

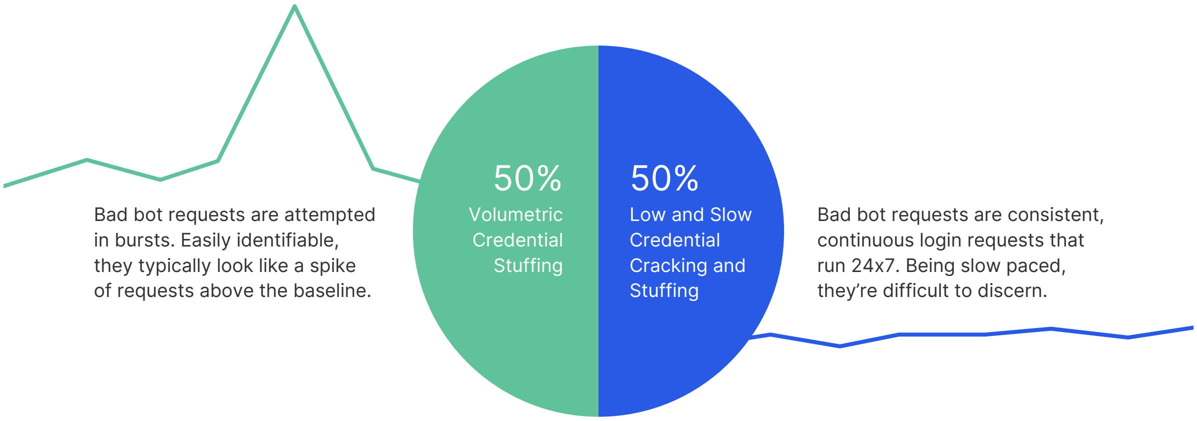

Bad bot requests are attempted in bursts. Easily identifiable, they typically look like a spike of requests above the baseline.

Bad bot requests are consistent, continuous login requests that run 24x7. Being slow paced, they're difficult to discern.

The Legacy of Data Breaches

Since 2013, the data breach reporting site haveIbeenpwned.com has reported an average of six breaches per month. They vary in size from a couple of thousand credentials to upwards of hundreds of millions (e.g., Onliner Spambot leaked 711,477,622 credentials in August, 2017). Meanwhile, data breach frequency has significantly increased during the last two years.

Every time there is a breach and credentials are made readily available, any business with a login page should get ready for a rise in volumetric credential stuffing attacks. Here, bot operators make two assumptions. The first is that people reuse their credentials on many websites. The second is that newly stolen credentials will be more likely to still be active. This is why businesses should anticipate bad bots running those credentials against their website after every breach.

With the increase in breaches and the billions of available stolen credentials, the rise of account and credential fraud activity by bot operators has been startling. Making matters worse, attackers are often well funded, highly skilled, motivated, and are unleashing sophisticated new techniques all the time to achieve their goals.

Volumetric Credential Stuffing Attacks by the Numbers

While the average number of volumetric attacks is 2 to 3 per month, some websites experience many more; one was hit by ten attacks in a single month. Volumetric attack distribution is shown below.

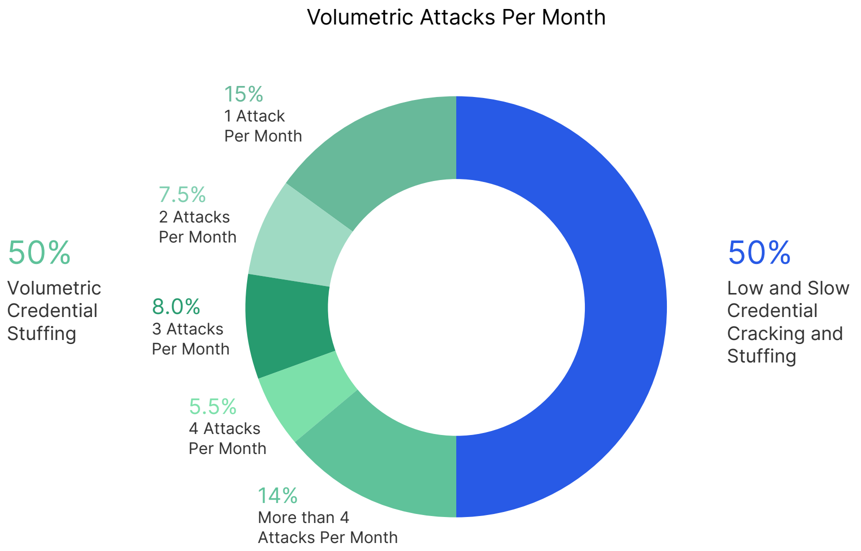

The Imperva network sees **17 volumetric credential stuffing attacks per day.**

2 - 3

NUMBER OF ATTACKS SEEN ON AVERAGE WEBSITE PER MONTH

3X

INCREASE IN VOLUMETRIC ATTACKS FOLLOWING A BREACH

10

MOST ATTACKS SEEN ON SINGLE WEBSITE IN ONE MONTH

Low-and-Slow Credential Cracking & Credential Stuffing Attacks

In contrast to targeted volumetric credential stuffing, where the attack is limited to a set timeframe and often appears as traffic spikes, the low-and-slow credential cracking and stuffing variety of attack are an ongoing, constant stream of malicious requests. Such attacks have no beginning or end, nor discernable pattern. They're more difficult to spot, as they test username and password combinations at a low-and-slow pace.

Low-and-slow attacks are similar to volumetric attacks in that they're highly distributed.

The Disruption of Account Takeover Attacks

Security teams know all about weekend disruptions, when alerts reveal that "more than fifty thousand login attempts have been made in only a matter of minutes." Someone from the team has to interrupt their weekend to address the problem, perhaps to discover the business has suffered an account takeover attack that began on Friday night.

Attacks Peak on Friday

When analyzing data across the global Imperva network, Friday and Saturday are when more attacks occur than the rest of the week. An explanation is that perhaps bot operators schedule attacks when it's presumed fewer people are around to notice anomalies. For security teams fed up with ruined weekends, installing a bot mitigation solution can take care of the problem.

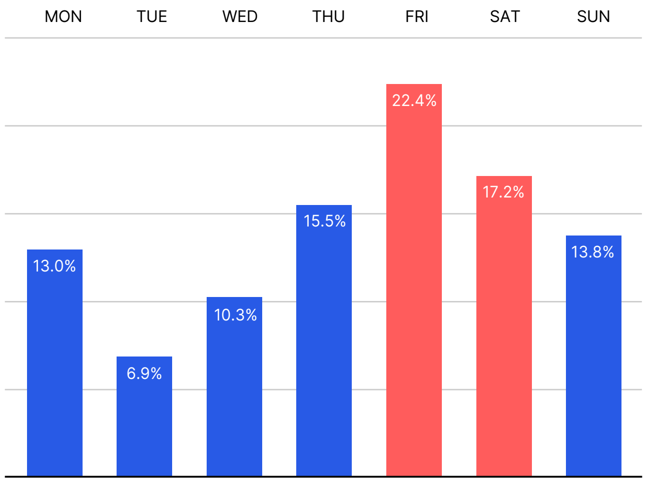

For security teams fed up with ruined weekends, installing a bot mitigation solution would take care of the problem.

You Can Set Your Clock to Bad Bots

We observe that attacks usually happen on a set frequency. For example, websites experiencing Wednesday attacks are likely to be hit by the next one on a Wednesday. Bot operators typically schedule attacks to launch automatically at the same time or day of the week.

Test Round Precedes Real Attack

Perpetrators tend to test their bad bots a few days before a large scale account takeover attack. Almost 20 percent of all analyzed attacks were preceded by a smaller scale test round a few days prior. While such tests are smaller in scale, any baseline anomaly from failed logins should be investigated and solutions considered.

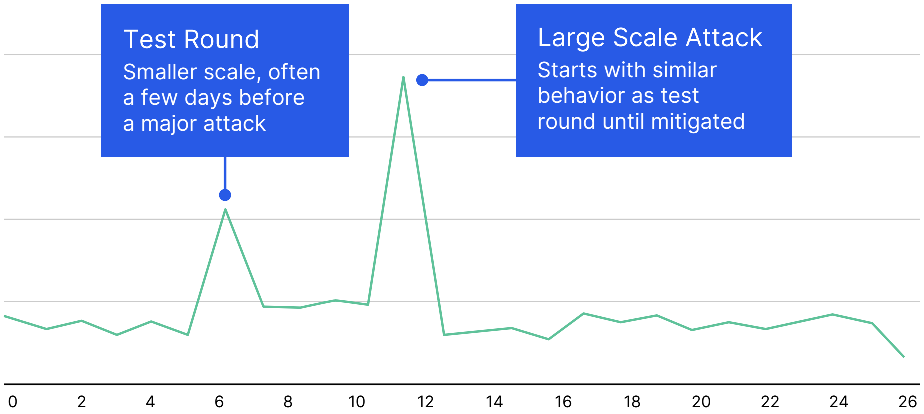

Test Round
Smaller scale, often a few days before a major attack

Large Scale Attack
Starts with similar behavior as test round until mitigated

20% test run their bots before a real attack.

Advanced and Evasive Behavior

The days of detecting attacks from a single source are past; sophisticated account takeover attacks are highly distributed. In large scale attacks, hundreds of thousands of IP addresses are used, each generating as few as two requests. Individually formatted and creating a unique profile, each is resubmitted from many different locations so as to appear legitimate.

Account Takeover Attack Profiles

The simple attack generates significant traffic per device fingerprint. It assumes that their many user agents will make them less suspicious. It is not distributed, making it easy to mitigate.

The moderate attack is much more distributed, but still generates a considerable amount of traffic per device fingerprint. The sophisticated attack shows as few as two requests per device fingerprint distributed over a large number of IP addresses.

	SIMPLE ATTACK	MODERATE ATTACK	SOPHISTICATED ATTACK
Total Requests	8,049	589,192	217,000
Login Attempts	8,020	589,192	205,102
Failed Login Increase	800%	370%	2,000%
IPs	14	20,753	17,491
User Agents	162	40	32
Device Fingerprints	2	37	72,105
IP Organizations	2	3,017	2,789
Countries	1	190	189

SIMPLE ATTACK: HOW TO BLOCK

Easily mitigate this attack by blocking using:

– IP address
– IP organizations

MODERATE ATTACK: HOW TO BLOCK

Mitigate by blocking by device fingerprint.

Blocking by IP address. IP organizations or countries is a never ending game of whack-a-mole and leads to high false positive rates.

SOPHISTICATED ATTACK: HOW TO BLOCK

Mitgate through deep interrogation to evalutate the legitimacy of each request.

By sending individual device configurations for almost every request blocking an identifier (IP address, country, IP organizations, fingerprint) is impossible.

THE SIZE OF VOLUMETRIC ACCOUNT TAKEOVER ATTACKS

35 - 50K

NUMBER OF REQUESTS IN AVERAGE CREDENTIAL STUFFING ATTACKS

500 - 5,000%

INCREASE IN TRAFFIC TO LOGIN PAGE

4 - 5 Million

NUMBER OF REQUESTS IN LARGE ATTACKS

Typical Post-Mitigation Behavior

Attacks tend to quickly die down once blocking is put in place. In most cases, bot operators try one or two different strategies before admitting defeat.

Persistent Sophisticated Account Takeover Attacks

When perpetrators are highly incentivized to obtain specific resource access and aren't simply validating credentials, they'll stick around and test new tactics. They continue until they're successful, or until continuing becomes more costly than their potential gains.

Bad Bot Operators Look for Other Targets

Whenever there is an easier attack target, bot operators use the path of least resistance. Attacking websites is generally easier than APIs or apps due to the effort required to blend in with normal traffic. But once blocked from a company's site, they're forced to find a new entry point—often a mobile app.

Common Account Takeover Tools

Sentry MBA

Sentry MBA is the most popular credential stuffing tool by far. Originally developed to test Sentry MBA's website for potential breaches, it has since been used by miscreants to execute large-scale credential stuffing attacks.

Sentry MBA lets an attacker configure it for any specific site, provide a proxy file to distribute requests, and the ability to use a list of credential combinations. It also includes built-in CAPTCHA solving. Many dark web forums and marketplaces offer configurations for specific web targets.

> DID YOU KNOW?
>
> Detect Sentry MBA in Logs: Sentry MBA enables changing of the user agent string, but many attackers still use the default configuration. It includes the following:
>
> - Opera/9.80 (Windows NT 6.0; U; en) Presto/2.2.0 Version/10.00
> - Mozilla/5.0 (Windows; U; Windows NT 5.1; en) AppleWebKit/522.11.3 (KHTML, like Gecko) Version/3.0 Safari/522.11.3
> - Mozilla/4.0 (compatible; MSIE 7.0; Windows NT 6.0; .NET CLR 1.1.4322; .NET CLR 2.0.50727; .NET CLR 3.0.4506.2152; .NET CLR 3.5.30729)
> - Mozilla/4.0 (compatible; MSIE 8.0; Windows NT 6.1; .NET CLR 1.1.4322; .NET CLR 2.0.50727; .NET CLR 3.0.4506.2152; .NET CLR 3.5.30729)
> - Mozilla/5.0 (Windows; U; Windows NT 5.1; en-US; rv:1.9.0.11) Gecko/ 2009060215 Firefox/3.0.11

Note that these user agent strings are used by others scripts and tools. But when observing them in combination, and within an observed credential stuffing incident, the likelihood of it being Sentry MBA is high.

When targeting mobile apps and APIs, bot operators often use tactics similar to those used on the website. Examples include assuming the target app identity or generating new identifiers in a highly distributed fashion.

Website to API and Back Again. Repeat.

Starting with spikes seen in the sample graph below, typical back and forth behavior is indicative of an attack on a website login page. Once a spike is mitigated on the website, a few attack attempts are made against the API. But then the perpetrator moves back to the web to have another go. This continues ad nauseum for many companies targeted by bad bots.

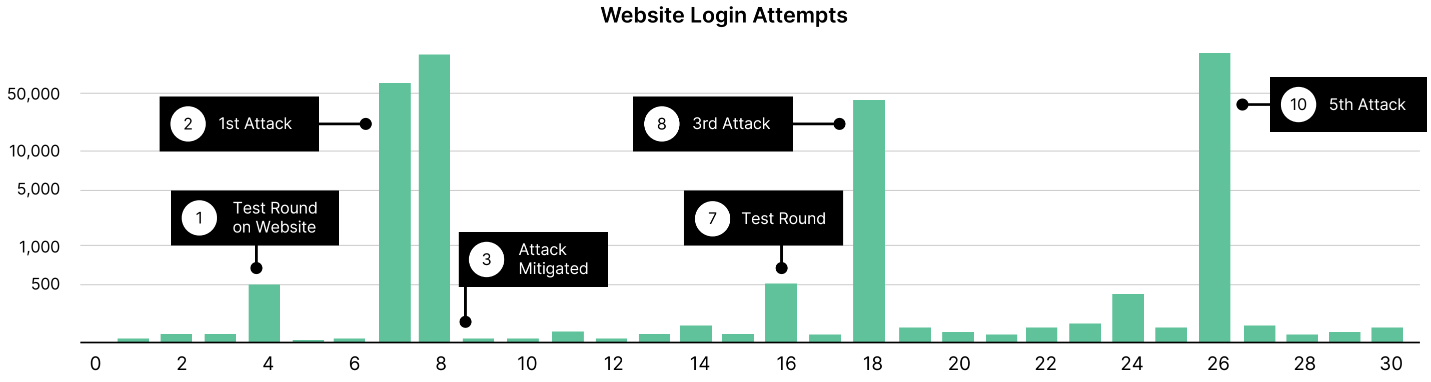

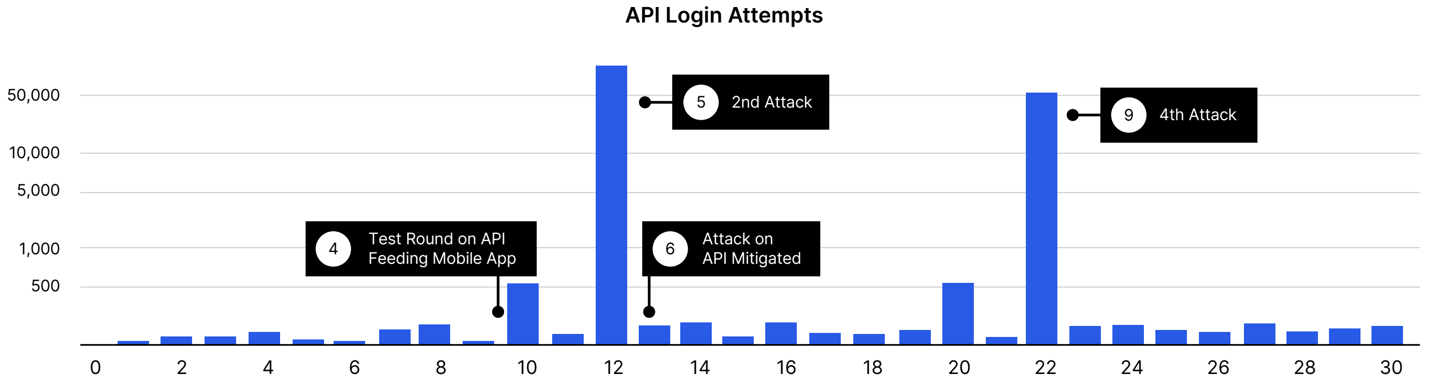

If blocking rules are shared between all systems, the risk of an attacker gaining access through another vector is less likely. But in all sophisticated account takeover attacks, bot mitigation must cover all access points— including web, API, and mobile apps.

Vertex

Released long before Sentry MBA (and rumoured to have been built by the same person) and being very similar, Vertex is still commonly used in attacks. It also uses a configuration, proxy, and credential file to execute attacks. And like Sentry MBA, Vertex can simultaneously use multiple, brute-force login interfaces.

> DID YOU KNOW?
>
> CAPTCHA Weakness – Vertex is easier to counter by using a CAPTCHA, as it doesn't have a built-in solving ability.

Account Hitmen

To use Account Hitman, one must have a sample HTTP request used during login, identify which response indicates success, upload a credentials file, then launch the tool against a target. Account Hitman distributes the requests via the provided proxy list to give it anonymity. Since the request format is based on users' original HTTP request, all formats will be valid (i.e., user agent, referrers, etc). Such user agents cannot be used to block requests, as they represent the most recent versions of all major browsers. To mitigate Selenium or other headless browsers, a specific bot defense mitigation needs to be in place.

> DID YOU KNOW?
>
> Detecting Account Hitman in Logs: Other than the provided requests, Account Hitman doesn't offer built-in functionality to alter headers, so all requests share the same format. An attack can be temporarily mitigated by applying rules around the observed request format. Account Hitman also doesn't have any solving abilities, so can be thwarted by CAPTCHAs.

Scripts & Headless Browsers (E.G., Selenium, Phantomjs)

On average, 30 percent of recent bad bot requests to login pages use full browser automation tools, such as Selenium. This represents around 5 percent of all login requests.

Selenium, PhantomJS, other headless browsers, and similar automation tools are still very popular among bot operators—credential stuffing attacks being no exception. Combining the power of these tools with simple scripts lets attackers quickly churn through stolen credentials lists.

> DID YOU KNOW?
>
> Link a Login Attack to Selenium: Compare user agents used in logs to the most common Selenium user agents seen during recent credential stuffing attacks:
>
> - Mozilla/5.0 (Windows NT 6.3; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/63.0.3239.132 Safari/537.36
> - Mozilla/5.0 (Windows NT 6.3; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/64.0.3282.167 Safari/537.36
> - Mozilla/5.0 (Macintosh; Intel Mac OS X 10_12_6) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/63.0.3239.132 Safari/537.36
> - Mozilla/5.0 (Windows NT 10.0; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/52.0.2743.116 Safari/537.36 Edge/15.15063
> - Mozilla/5.0 (Windows NT 6.3; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/64.0.3282.186 Safari/537.36

Such user agents cannot be used to block requests, as they represent the most recent versions of all major browsers. To mitigate Selenium or other headless browsers, a specific bot defense mitigation needs to be in place.

Recommendations for Detecting Bad Bot Activity

Bots are on your website every day, and attack characteristics become more advanced and very nuanced. How should businesses go about protecting themselves? Every site is targeted for different reasons, and usually by different methods, so there is no one-size-fits-all bot defense solution. But there are some proactive steps you can take to start addressing the problem.

1. **BLOCK OR CAPTCHA OUTDATED USER AGENTS/BROWSERS:** The default configurations for many tools and scripts contain user-agent string lists that are largely outdated. This step won't stop the more advanced attackers, but it might catch and discourage some. The risk in blocking outdated user agents/browsers is very low; most modern browsers force auto-updates on users, making it more difficult to surf the web using an outdated version. We recommend you block or CAPTCHA the following browser versions:

	BLOCK End of Life More than 3 years	CAPTCHA End of Life More than 2 years
Firefox version	< 38	< 45
Chrome version	< 41	< 49
Internet Explorer version	< 10	10
Safari version	< 9	9

2. **BLOCK KNOWN HOSTING PROVIDERS AND PROXY SERVICES:** Even if the most advanced attackers move to other, more difficult-to-block networks, many less sophisticated perpetrators use easily accessible hosting and proxy services. Disallowing access from these sources might discourage attackers from coming after your site, API, and mobile apps.

   BLOCK THESE DATA CENTERS:

   - DigitalOcean
   - OVH SAS
   - Choopa, LLC
   - OVH Hosting
   - GigeNET
   - Amazon.com

3. **PROTECT EVERY BAD BOT ACCESS POINT:** Be sure to protect exposed APIs and mobile apps—not just your website—and share blocking information between systems wherever possible. Protecting your website does little good if backdoor paths remain open.

4. **CAREFULLY EVALUATE TRAFFIC SOURCES:** Monitor traffic sources carefully. Do any have high bounce rates? Do you see lower conversion rates from certain traffic sources? These can be signs of bot traffic.

5. **INVESTIGATE TRAFFIC SPIKES:** Traffic spikes appear to be a great win for your business. But can you find a clear, specific source for the spike? One that is unexplained can be a sign of bad bot activity.

6. **MONITOR FOR FAILED LOGIN ATTEMPTS:** Define your failed login attempt baseline, then monitor for anomalies or spikes. Set up alerts so you're automatically notified if any occur. Advanced low-and-slow attacks don't trigger user or session-level alerts, so be sure to set global thresholds.

7. **MONITOR INCREASES IN FAILED VALIDATION OF GIFT CARD NUMBERS:** An increase in failures, or even traffic, to gift card validation pages can be a signal that bots such as GiftGhostBot are attempting to steal gift card balances.

8. **PAY CLOSE ATTENTION TO PUBLIC DATA BREACHES:** Newly stolen credentials are more likely to still be active. When large breaches occur anywhere, expect bad bots to run those credentials against your site with increased frequency.

9. **EVALUATE A BOT MITIGATION SOLUTION:** The bot problem is an arms race. Bad actors are working hard every day to attack websites across the globe. The tools used constantly evolve, traffic patterns and sources shift, and advanced bots can even mimic human behavior. Hackers using bots to target your site are distributed around the world, and their incentives are high. In early bot attack days you could protect your site with a few tweaks; this report shows that those days are long gone. Today it's almost impossible to keep up with all of the threats on your own.

Imperva is an analyst-recognized, **cybersecurity leader** championing the fight to **secure data and applications** wherever they reside.