radware

Intent-based Deep Behavioral Analysis:
A Proprietary Bot Detection Technology

WHITE PAPER

TABLE OF CONTENTS

⊖ Introduction

Over half of all internet traffic is generated by bots — some legitimate, some malicious. Competitors and adversaries alike deploy "bad" bots that leverage different methods to achieve nefarious objectives. This includes account takeover, scraping data, denying available inventory and launching denial-of-service attacks with the intent of stealing data or causing service disruptions.

These attacks often go undetected by conventional mitigation systems and strategies because bots have evolved from basic scripts to large-scale distributed bots with humanlike interaction capabilities to evade detection mechanisms. To stay ahead of the threat landscape requires more sophisticated, advanced capabilities to accurately detect and mitigate these threats. In this white paper, we discuss one of the key technical capabilities required to stop today's most advanced bots: intent-based deep behavioral analysis (IDBA).

IDBA is a major step forward in bot detection technology because it performs behavioral analysis at a higher level of abstraction of intent, unlike the commonly used, shallow interaction-based behavioral analysis. For example, account takeover is an example of an intent, while "mouse pointer moving in a straight line" is an example of an interaction. Capturing intent enables IDBA to provide significantly higher levels of accuracy to detect advanced bots. IDBA is designed to leverage the latest developments in deep learning.

In this white paper, we first present a brief overview of the evolution of bots and the corresponding detection technologies, in addition to identifying major challenges in applying behavioral analysis to detect advanced bots. Finally, we describe how IDBA works and how it addresses these challenges.

Key Features

IDBA uses semisupervised learning models to overcome the challenges of inaccurately labeled data, bot mutation and the anomalous behavior of human users.

IDBA leverages **intent encoding, intent analysis and adaptive-learning** techniques to accurately detect large-scale distributed bots with sophisticated humanlike interaction capabilities.

Evolution of Bots and Corresponding Detection Technologies

Generally speaking, bots can be categorized into four categories, or levels, based on their degree of sophistication (see Figure 1).
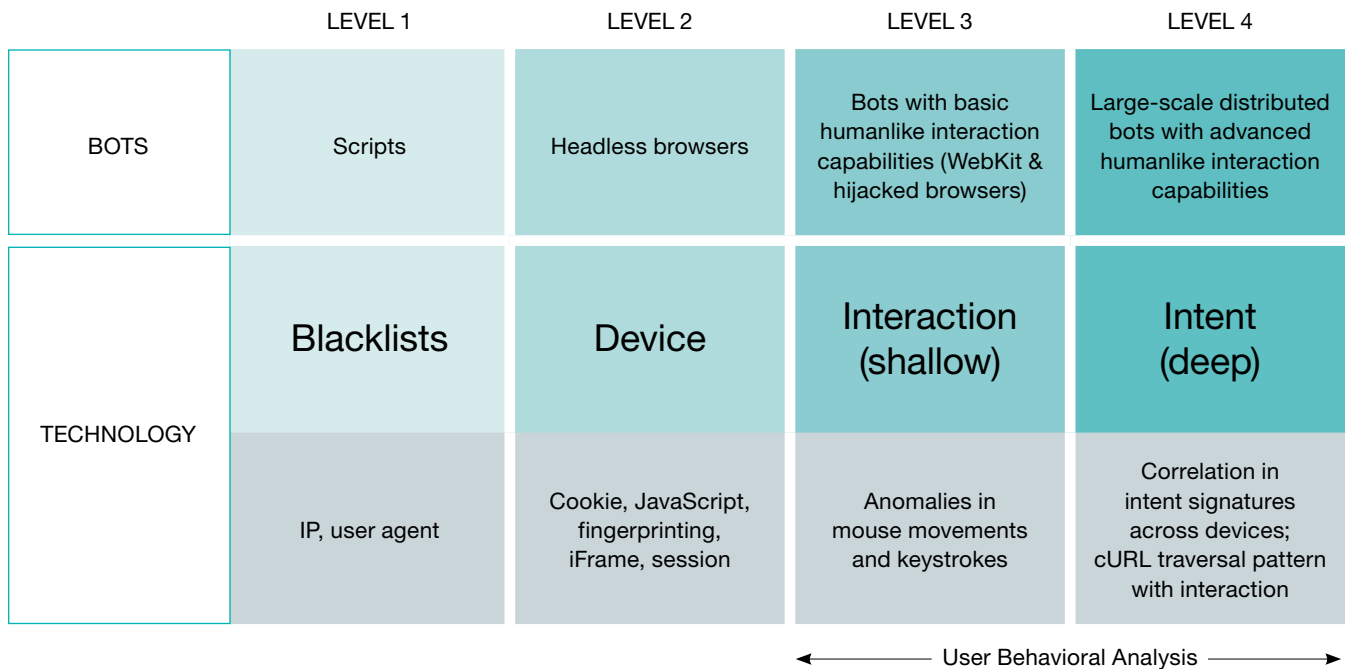
	LEVEL 1	LEVEL 2	LEVEL 3	LEVEL 4
BOTS	Scripts	Headless browsers	Bots with basic humanlike interaction capabilities (WebKit & hijacked browsers)	Large-scale distributed bots with advanced humanlike interaction capabilities
TECHNOLOGY	Blacklists	Device	Interaction (shallow)	Intent (deep)
	IP, user agent	Cookie, JavaScript, fingerprinting, iFrame, session	Anomalies in mouse movements and keystrokes	Correlation in intent signatures across devices; cURL traversal pattern with interaction

←————— User Behavioral Analysis —————→

Figure 1: The four levels of bots and the capabilities required to mitigate each one successfully

Level 1 bots: These are built using basic scripting and make cURL-like requests to web properties from a small number of IP addresses. Such bots can be identified through a blacklist of user agents (UAs) as well as a basic analysis of the traffic emanating from the IP address. These bots typically lack the capability to store cookies or to execute JavaScript, which can be identified through a simple analysis of the traffic. The bot can then be mitigated through blacklisting of its IP address and UA combination.

Level 2 bots: These bots operate through headless browsers, such as PhantomJS, and are able to store cookies and execute JavaScript. However, most of them can be identified through their browser and device characteristics, such as the presence of specific JavaScript variables, iFrame tampering, sessions, cookies, etc. Once the bot is identified, it can be blocked using its fingerprint.

Level 3 bots: These bots use full-fledged browsers (dedicated or hijacked) for their operations. They are also able to simulate basic humanlike interactions, such as simple mouse movements and keystrokes. This makes it difficult to detect these bots based on device and browser characteristics. Interaction-based user behavioral analysis can be leveraged to detect such bots. For example, specific interaction characteristics, such as "movement of mouse pointer in a straight line" or "very rapid keystrokes," can be analyzed to detect such bots.

Level 4 bots: These bots have advanced humanlike interaction characteristics, such as moving the mouse pointer in random humanlike patterns instead of in straight lines. Detecting these bots based on shallow interaction characteristics like mouse movement patterns will result in a high number of false positives. These bots are also massively distributed across tens of thousands of IP addresses. Prevailing techniques are therefore inadequate in mitigating such bots. Radware's Bot Manager technology performs intent-based deep behavioral analysis to detect such bots with high accuracy.

→ Challenges of Using Behavioral Analysis to Detect Bots

As mentioned earlier, shallow interaction-level behavioral analysis is insufficient to detect advanced Level 4 bots. This points to the need to perform behavioral analysis at a deeper level: The visitor's journey through the web property needs to be analyzed in addition to the interaction-level characteristics, such as mouse movements. This includes additional signals such as the sequence of URLs traversed, the referrers used and the time spent on each page. Using this richer behavioral information, the incoming visitor can be classified as a human or bot.

The complexity and diversity of signals such as interaction and URL traversal information naturally point to the application of machine learning — in particular, deep learning — to perform this classification task. However, to apply machine learning for this task, a series of challenges must first be addressed.

▷ **Garbage in, garbage out (GIGO):** The recent successes of deep learning have been in the domain of supervised learning, in which labeled data (i.e., data on past human and bot traffic) is available to train the machine learning model. However, direct application of supervised learning to bot detection is problematic.

Level 3 bot detection techniques are insufficient to detect Level 4 bots accurately. This means that currently available data on Level 4 bots in a bot detection system will not be correctly labeled. There will be a significant number of false positives and/or false negatives due to misclassification. A supervised learning model trained using this data will have the same issues as the Level 3 bot detection techniques since the model's performance is directly related to the quality of the available data. Hence addressing the GIGO challenge is a central problem to the application of machine learning in the security domain.

▷ **Bot mutation:** When a bot gets detected and blocked, it changes its characteristics to improve its evasive capabilities. A basic supervised learning model trained by using past data will have difficulty in detecting such mutations.

▷ **Anomalous characteristics of human visitors:** Given the above limitations of supervised learning, unsupervised learning is another approach to be considered. Unsupervised learning identifies hidden patterns in unlabeled data such as grouping or anomalies. As it does not rely on labels, it is not affected by the issues posed by GIGO and bot mutation. Unsupervised learning can help in identifying bots with anomalous characteristics (anomaly detection) and bot clusters (clustering). However, certain human visitors can also have anomalous characteristics or grouping. For example, certain users of a web property may be power users with higher-than-average levels of engagement. These users could get flagged as anomalies or clusters. Thus, a straightforward application of unsupervised learning to bot detection can result in false positives.

Radware's Bot Manager runs with IDBA technology to overcome the challenges outlined above. IDBA technology addresses these challenges using a semisupervised approach coupled with adaptive learning.

The Three Stages of IDBA

1. **Intent encoding:** The visitor's journey through the web property is captured through signals such as mouse or keystroke interactions, URL and referrer traversals, and time stamps. These signals are encoded using a proprietary, deep neural network architecture into an intent encoding-based, fixed-length representation. The encoding network jointly achieves two objectives:

   - To be able to represent the anomalous characteristics of completely new categories of bots

   - To provide greater weight to behavioral characteristics that differ between humans and bots

2. **Intent analysis:** Here, the intent encoding of the user is analyzed using multiple machine learning modules in parallel. A combination of supervised and unsupervised learning-based modules are used to detect both known and unknown patterns.

3. **Adaptive learning:** The adaptive-learning module collects the predictions made by the different models and takes actions on bots based on these predictions. In many cases, the action involves presenting a challenge to the visitor like a CAPTCHA or an SMS OTP that provides a feedback mechanism (i.e., CAPTCHA solved). This feedback is incorporated to improvise the decision-making process. Decisions to be made can be broadly categorized into two types of tasks.

   - **Determining thresholds:** The thresholds to be chosen for anomaly scores and classification probabilities are determined through adaptive threshold control techniques.

   - **Identifying bot clusters:** Selective incremental blacklisting is performed on suspicious clusters. The suspicion scores associated with the clusters (obtained from the collusion detector module) are used to set prior bias.

→ Conclusion

Current bot detection and classification methodologies are ineffective in countering the threats posed by rapidly evolving and mutating sophisticated bots. Bot detection techniques that use interaction-based behavioral analysis can identify Level 3 bots but fail to detect the advanced Level 4 bots that have humanlike interaction capabilities. The unavailability of correctly labeled data for Level 4 bots, bot mutations and the anomalous behavior of human visitors from disparate industry domains required the development of semisupervised models that work at a higher level of abstraction of intent, unlike only interaction-based behavioral analysis. Radware's Bot Manager IDBA overcomes the challenges faced by straightforward application of supervised and unsupervised machine learning models. IDBA leverages a combination of intent encoding, intent analysis and adaptive-learning techniques to identify the intent behind attacks perpetrated by massively distributed humanlike bots.

About Radware