SANS

Analyst Program

Product Review

SANS Review: reCAPTCHA Enterprise

Written by **Dave Shackleford**

July 2023

Google Cloud

Introduction

With more web-based services and online business being conducted than ever before, it's no surprise that attacks and fraud activity have escalated significantly in the past few years. PwC's Global Economic Crime and Fraud Survey 2022 found that technology companies saw a 53% increase in attempted fraud compared to a year prior, with many losing more than $1 million as a result.[1]

To combat online fraud, malicious bots, and account takeover, organizations have traditionally turned to the CAPTCHA (Completely Automated Public Turing test to tell Computers and Humans Apart) mechanism to discern between real human beings interacting with sites and application services versus automated bots with no contextual behavior capabilities. To many, CAPTCHAs have been viewed as a necessary evil that creates a "speed bump" of sorts for certain types of site interaction, often asking users to identify pictures with certain types of images present. Over time, though, the back-end analysis of bot interactions has improved, ideally making CAPTCHA controls more integrated and automated.

SANS had the opportunity to review the Google reCAPTCHA Enterprise platform. reCAPTCHA Enterprise employs advanced machine learning algorithms to analyze user behavior and assess risk levels in real time. By identifying and blocking automated bots and fraud actors, reCAPTCHA Enterprise helps secure online services against account takeovers, unauthorized access, and fraudulent activities. reCAPTCHA Enterprise boasts some of the following capabilities and advantages:

- Enhanced risk scoring (11 scores between 0 and 1)
- Frictionless risk assessment no longer requiring visual challenges for end users and risk scores returned to developers directly
- Available API integration and "reason codes" that can be passed to application and security logic models and controls
- Integrated risk-based authentication engine (Account Defender), multifactor authentication (MFA), and password leak detection for a full account takeover and bulk fake-account protection
- Integrated fraud-prevention engine to thwart fraudulent payment transactions and reduce erroneous rejections
- Native mobile software development kits (SDKs) for Android and iOS
- Native integration with Google's Cloud Armor web application firewall (WAF), making Google Cloud application deployment more secure and seamless, protecting customers at the network edge

[1] www.pwc.com/gx/en/services/forensics/economic-crime-survey.html

Other key features of reCAPTCHA Enterprise include:

- **Adaptive security—**reCAPTCHA Enterprise adapts to new threats and evolves with your security needs, ensuring a robust defense against ever-changing attack vectors.
- **Customizable security rules—**Tailor your security settings based on your organization's risk tolerance and specific use cases.
- **Frictionless user experience—**Minimize user friction by maintaining a smooth and secure browsing experience.

reCAPTCHA Enterprise's fraud-prevention engine is available through web services, native mobile SDKs (iOS/Android), and WAF integrations (Cloud Armor and Fastly currently). reCAPTCHA Enterprise provides a comprehensive online fraud-detection platform that helps prevent fraudulent and spammy or abusive digital client activity across any web and application services where deployed. The complete platform includes Account Defender, which protects users from account takeovers (ATOs); frictionless bot management capabilities; and password leak detection to identify compromised accounts. With an additional API call for fraud prevention, organizations can leverage additional solutions to secure financial transactions along with core reCAPTCHA Enterprise capabilities. This whitepaper includes additional information about the platform, although we didn't observe or evaluate some features.

Password Leak Detection[2]

In recent years, there have been a plethora of leaked accounts and credentials on the internet related to data breaches. Malicious actors often purchase or acquire these credentials to use them in wide-ranging attacks against authentication pages in a phenomenon known as *credential stuffing*. Credential stuffing attacks leverage leaked or stolen credentials to gain unauthorized access to user accounts. reCAPTCHA Enterprise's password leak detection capability helps protect your users by:

- **Monitoring the Dark Web—**Google's scanning engines and threat-hunting teams search for leaked credentials on the Dark Web, alerting you to potential security breaches.
- **Detecting password leaks—**When you enable password leak detection, any attempts to log in to applications and services with known leaked credentials can immediately alert on a detected password leak and also trigger a password reset to protect affected users.
- **Enhanced login security—**If desired, organizations can enable an extra layer of security (step-up authentication) to the login process, mitigating the risk of account takeovers.

One of the biggest concerns with credential-centric threat intelligence and monitoring is privacy and protection of the credential data. Google password leak detection uses an encryption technique called *multiparty computation* to protect users' credentials. This technique means that reCAPTCHA never sees the actual credentials in plaintext, so they can't be exposed or leaked in any way. Through these privacy-preserving measures, password leak detection offers a more secure and responsible solution for monitoring leaked credentials, helping protect users from account takeovers while maintaining privacy and trust.

[2] https://cloud.google.com/recaptcha-enterprise/docs/check-passwords

Web Risk Submission API[3]

In alignment with other key areas of threat intelligence and online monitoring, the Google Web Risk Submission API can help to mitigate phishing attacks that pose a significant threat to organizations and their brand reputation. The Google Web Risk Submission API provides a powerful tool to help safeguard your brand by identifying and blocking phishing pages that target your organization from more than 5 billion devices. By leveraging this API, you can submit potential phishing URLs to Google, which will then be analyzed and, if confirmed as malicious, added to the Google Safe Browsing list in minutes.

Many web browsers take advantage of this list to warn users of potentially harmful websites, thus preventing any visitors to applications and sites from falling victim to phishing attacks. Implementing the Web Risk Submission API not only bolsters defenses against phishing threats, but also fosters trust among application users by demonstrating commitment to online safety. Given the alignment between online credential hijacking and phishing, integrating reCAPTCHA with phishing monitoring and prevention makes sense. When paired with password leak detection, this helps provide a full credential-loss protection tool for organizations to keep their users safe.

Account Defender[4]

One of the most significant attack vectors we face today is account hijacking. Adversaries have realized that end users (and their application and domain accounts) are much easier targets than exposed services in many cases, and attackers often initiate ingress during campaigns by leveraging compromised user accounts. The reCAPTCHA Enterprise Account Defender capability helps safeguard online systems and applications by detecting and preventing account takeover and bulk fake-account creation. By leveraging machine learning and advanced risk analysis, Account Defender helps to:

- **Block credential stuffing attacks and bulk account creation—**Account Defender can help to identify patterns of suspicious activity on login and registration pages to help protect users against account takeovers and abuse.
- **Reduce friction on login for legitimate users—**By identifying "normal" users and logins, Account Defender can help to simplify legitimate user logins and increase friction for suspicious users on login.
- **Reduce SMS costs—**Although much has been debated regarding SMS as a multifactor authentication method, many organizations and "step-up authentication" controls models still employ this as a means of out-of-band user validation and verification. With additional user account threat intelligence, Account Defender can help to minimize costs associated with SMS verification by blocking fake accounts before they are registered in applications and by only challenging users on risky logins. Profile matches can be let through without a challenge, thus reducing the cost of supporting SMS as a second-factor authentication.

[3] https://cloud.google.com/web-risk/docs/submission-api

[4] https://cloud.google.com/recaptcha-enterprise/docs/account-defender

By integrating the Related Accounts API[5] with Account Defender, organizations can also:

- **Uncover hidden connections—**One of the most challenging issues for security operations is mapping and monitoring end user behaviors at scale. Account Defender can help to detect and analyze account relationships, allowing security teams to identify unusual or suspicious patterns of behavior, such as multiple fake accounts linked to a single malicious actor.
- **Strengthen security measures—**Organizations can apply security rules to groups of related accounts, ensuring that a single compromised account does not lead to a flood of account takeovers within the network environment.
- **Optimize fraud detection—**Machine learning and strong back-end analytics enhance the accuracy of Account Defender's risk analysis by factoring in the relationships between accounts, enabling more precise detection of bots and fraudulent activities.

Using the Related Accounts API along with Account Defender not only bolsters security against bulk fake account creation, but also provides valuable insights into the broader context of user activity.

Reviewing reCAPTCHA Enterprise

For our review, we actually deployed a sample site within a Google Cloud account to evaluate the various use cases that reCAPTCHA Enterprise highlights. To get started in a typical deployment, organizations can install a score-based site key on each part of the application front end and send transaction data when an interaction occurs. This helps train site-specific fraud models and starts returning fraud scores for each transaction.

To kick this off, we simply accessed the reCAPTCHA Enterprise category within the Security section of our Google Cloud Platform (GCP) account and selected Try It Out (see Figure 1).

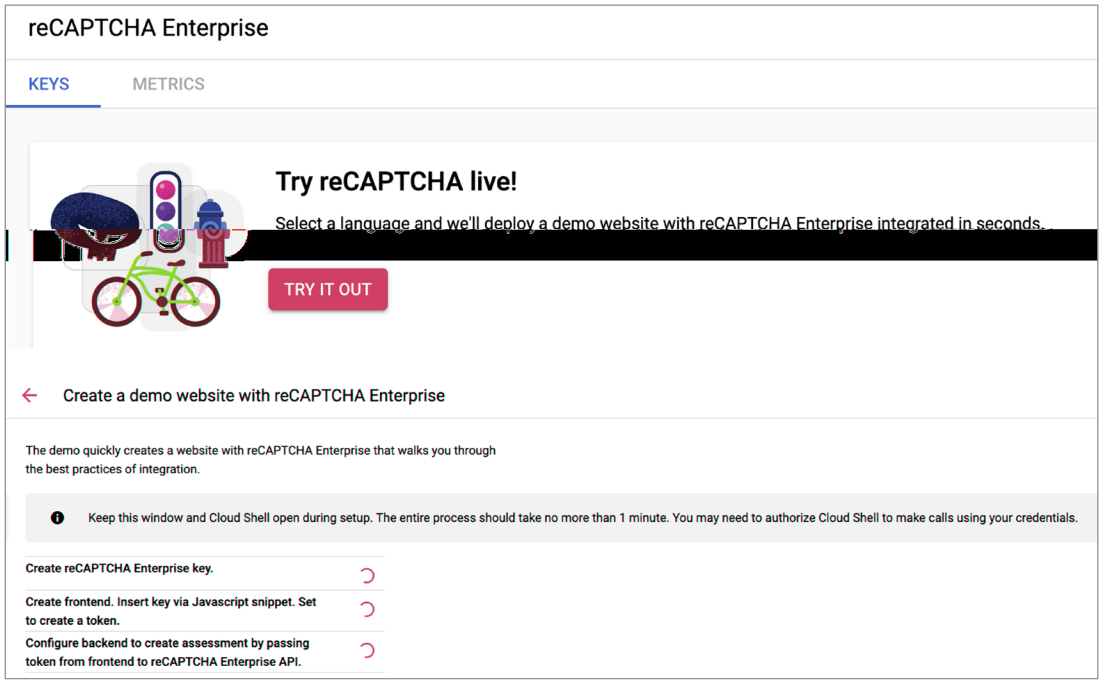

Figure 1. Initiating the reCAPTCHA Enterprise Demo Site

[5] https://cloud.google.com/recaptcha-enterprise/docs/account-query-apis

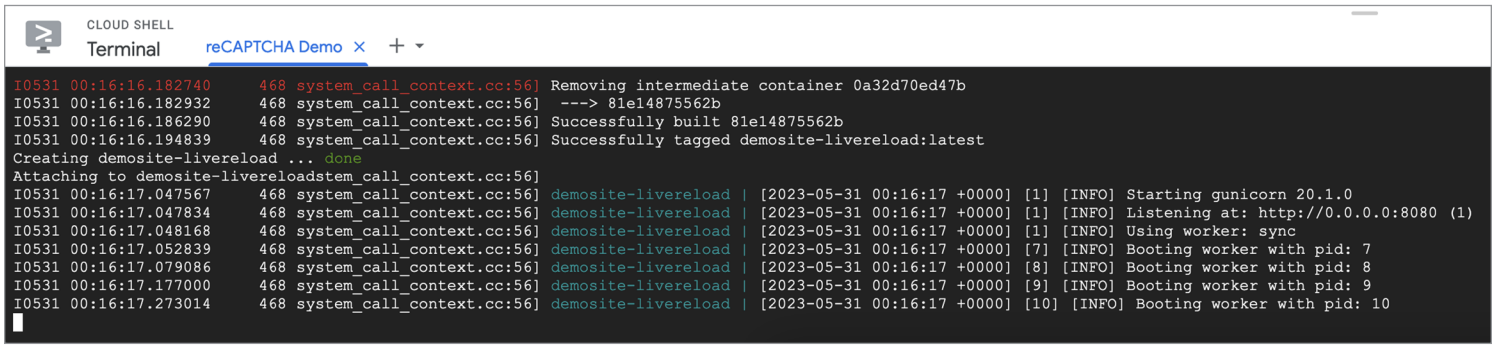

Figure 2. reCAPTCHA Enterprise Site Initiating

Next, we waited while the various back-end containers and site components were started. Google makes this very simple through the Google Cloud Shell, which you simply have to authorize (see Figure 2).

After the site has finished generating, you can modify the web application pages as desired and manage the status of the back-end services and APIs as well. In this case, we simply initiated the default site (see Figure 3).

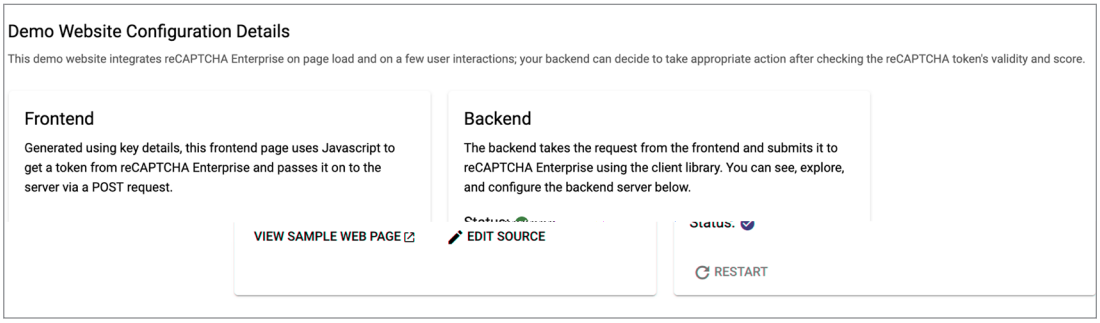

Figure 3. reCAPTCHA Enterprise Default Site Initiation

Once the site was up and running, we chose to view the sample page and started the walkthrough of the reCAPTCHA Enterprise use cases.

Protecting Access to an Entire Website

The first, and perhaps most simple, use case with reCAPTCHA Enterprise is protecting the initial access to a website. The most common scenario where additional validation is required for access in this case is known bots or malicious/suspicious sites and domains attempting access, or repeated patterns of access that indicate malicious behavior. As Figure 4 shows, the initial site for our reCAPTCHA Enterprise review is a sample game called BadFinder.

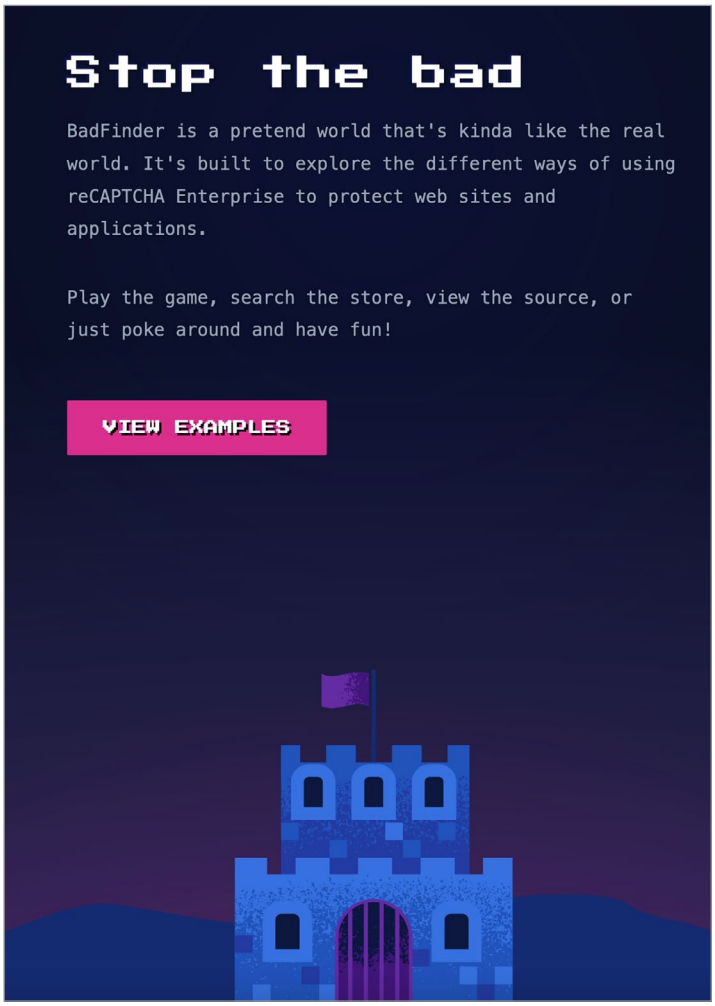

Figure 4. The reCAPTCHA Enterprise Sample Site

Because we have not accessed the site before and don't exhibit any malicious or suspicious behaviors, Google ranks the access attempt at 90% confident of legitimate behavior, as shown in Figure 5.

Additional verification tokens can be integrated for valid users if desired, increasing the behavioral analysis score for access to protected sites.

User Account Creation or Signup

Another common bot attack technique is to abuse account creation pages to generate fake or "synthetic" accounts that can be used to maliciously interact with applications and data or target other online resources. Bots often try to misuse an application's account signup to impact and simulate "real" customer experiences as they prepare for larger-scale attacks. Common signs of unusual account creation may include:

- Higher-than-average account creation rate
- Accounts with incomplete information
- Accounts created and not used immediately
- Accounts created and used inappropriately relative to their assigned roles

We browsed to the reCAPTCHA Enterprise signup page (see Figure 6).

Figure 5. A Normal reCAPTCHA Enterprise Interaction

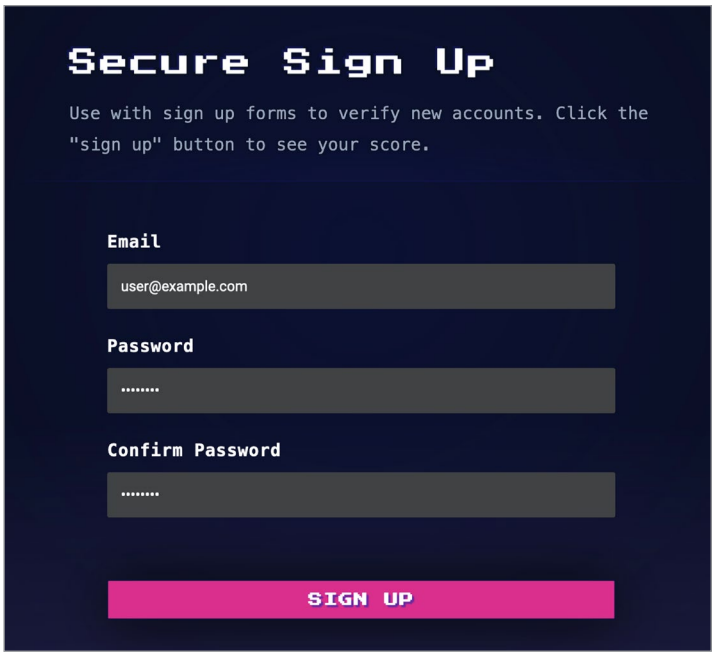

Figure 6. reCAPTCHA Enterprise Demo Account Creation Page

In this case, because we manually interacted with the signup form fields and entered information, coming from a benign source address/domain, we were viewed as not being malicious by reCAPTCHA Enterprise (see Figure 7).

By analyzing the frequency and types of account creation occurring, reCAPTCHA Enterprise can help to identify unusual patterns of account generation activity. By pairing this with services like Account Defender and password leak detection to look for known compromised accounts, leaked credentials, and other threat intelligence related to accounts, organizations can easily cut down on possible malicious account creation at scale. When paired with the offline Related Accounts API, this enables the detection of related accounts associated with the same attacker, thus allowing for scalable detection.

Malicious Logins

Another related type of bot activity (as well as targeted campaigns with malicious actors) relates to obvious malicious logins. We browsed to the reCAPTCHA Enterprise login page to assess how it viewed our attempted login (see Figure 8).

Figure 7. Account Creation Risk Scoring

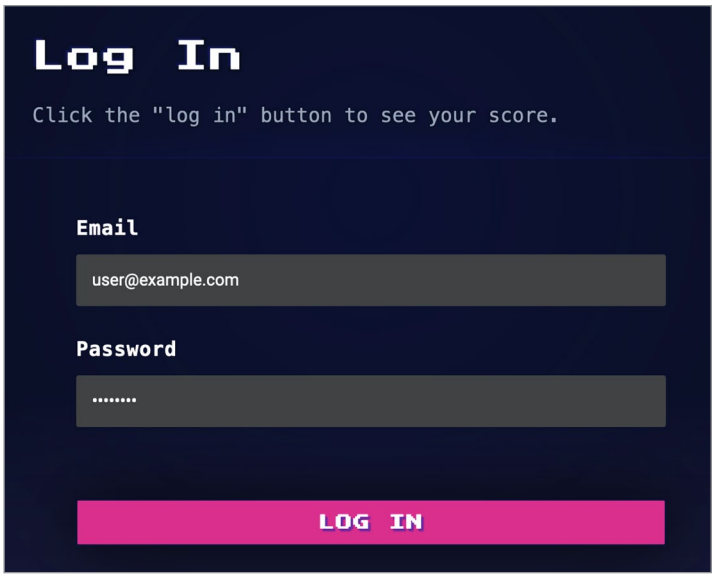

Figure 8. The reCAPTCHA Enterprise Login Page

As with the previous example (creating a new user account), our interaction with the site was seen as normal and unsuspicious (see Figure 9).

The tie-in with Account Defender can help to demonstrate how reCAPTCHA also returns bot risk scores. This enables developers to determine whether the attack is automated. Account Defender answers this question: Is the right user coming back this time? If so (due to matching attributes like originating IP address, end user equipment, browser, geographic location, and hashed user account, for example), reCAPTCHA can recommend skipping additional authentication steps via MFA, streamlining the user experience (see Figure 10).

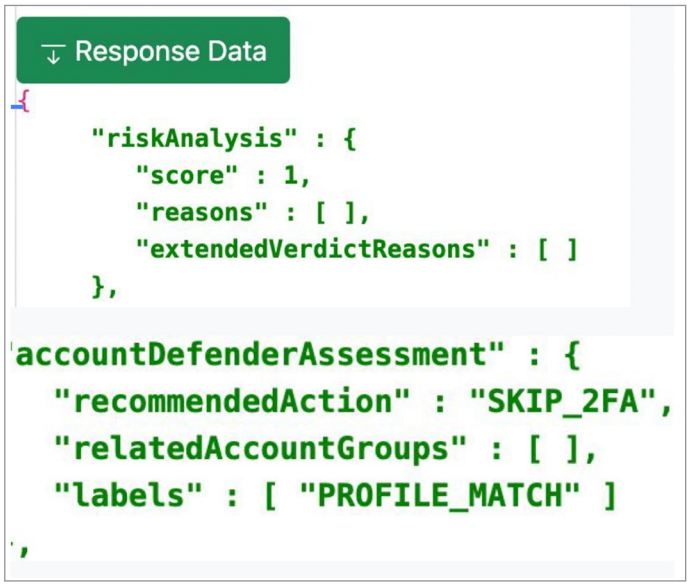

Figure 10. A Trusted User Interaction

For risky connections (those that come from different locations or systems that exhibit unusual connection patterns, for example), MFA could be requested (see Figure 11).

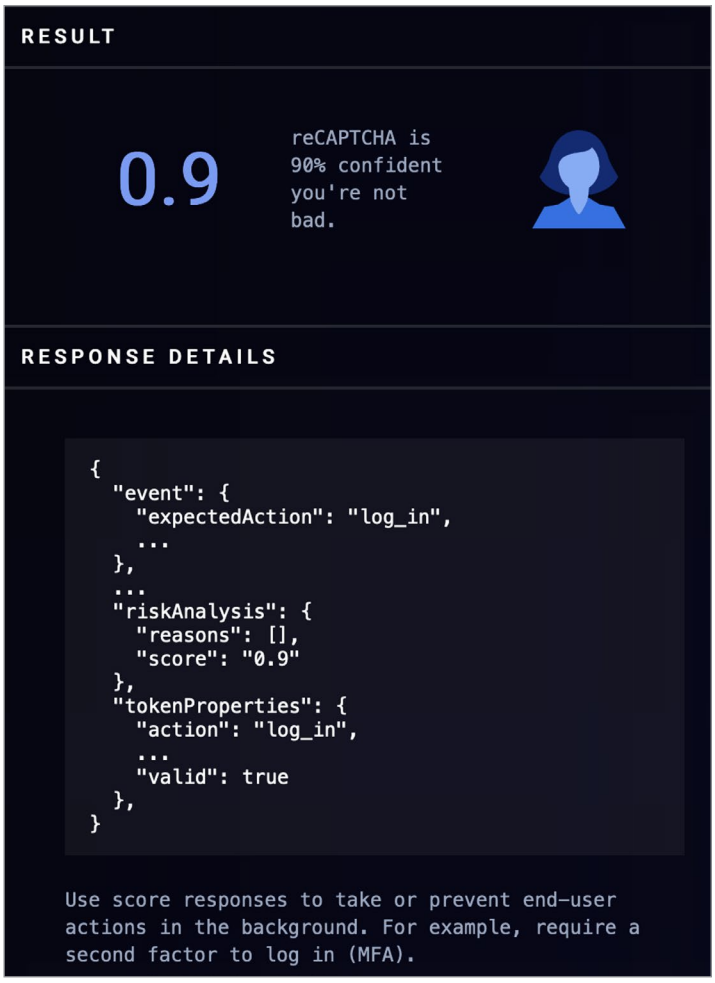

Figure 9. A Safe/Normal Login Interaction with reCAPTCHA Enterprise

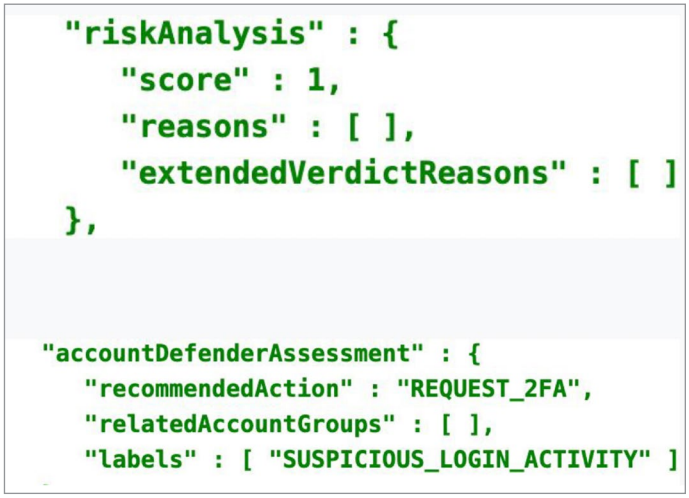

Figure 11. More Advanced Authentication and Access Analysis

In this example, the login comes back as "SUSPICIOUS_LOGIN_ACTIVITY", meaning this account is at risk of an account takeover. In this case, Google recommends that MFA be required ("REQUEST_2FA"). reCAPTCHA can also easily identify and label failed and successful MFA interactions (see Figure 12).


```
{
    "riskAnalysis" : {
       "score" : 1,
       "reasons" : [ ],
       "extendedVerdictReasons" : [ ]
    },
    "accountVerification" : {
       "endpoints" : [ ],
       "latestVerificationResult" : "SUCCESS_USER_VERIFIED",
       "languageCode" : "",
       "username" : ""
    },
```


Figure 12. Successful MFA Authentication and User Verification

Another variation of unusual activity could be related to attempts to create new user accounts. Account Defender analyzes how the account was created, possibly indicating that the user account was created maliciously or suspiciously, likely by bots. As shown in Figure 13, the return of the "SUSPICIOUS_ACCOUNT_CREATION" label prompts developers and security teams to scrutinize the account more carefully, to consider adding additional validation and authentication requirements to the account, and to ensure they note no other suspicious activity.


```
"accountDefenderAssessment" : {
   "recommendedAction" : "RECOMMENDED_ACTION_UNSPECIFIED",
   "relatedAccountGroups" : [ ],
   "labels" : [ "SUSPICIOUS_ACCOUNT_CREATION" ]
},
```


Figure 13. Suspicious Account Creation

Online Purchasing

One of the most prevalent types of online fraud relates to online purchases in e-commerce. Malicious actors employ a number of common attack vectors and methods when committing payment fraud:

- **Carding—**When attackers acquire stolen payment card data, they may repeatedly try to validate stolen card numbers and credentials through storefronts to see if they are still active.
- **Card cracking—**Similar to the carding method, card cracking is a brute-force attempt to validate payment card security codes on e-commerce sites.
- **Cashing out—**Cashing out is a simple scam that focuses on purchases with stolen payment card data or credit that can be exchanged for "chargeback" money in returns.
- **Coupon fraud—**By identifying valid payment token codes, attackers may be able to gain access to valid cash alternatives and limited offers for legitimate consumers, which can impact customer loyalty programs.

reCAPTCHA Enterprise fraud prevention helps protect payment transactions by identifying targeted manual attacks and large-scale fraud attempts. It automatically trains fraud models based on behavior and transaction data to identify events that are likely fraudulent and could cause a dispute or chargeback if accepted. When threat actors are identified and blocked on any site in the reCAPTCHA Enterprise network, the intelligence is made available to help protect other organizations from those same attackers. Organizations can then use the scores to either manually review the transaction or block suspicious transactions outright. This helps to increase trust in their legitimate transactions, reduce the amount of friction for genuine users, and reduce erroneous rejection rates.

We evaluated the perceived risk level of our interactions by visiting the reCAPTCHA Enterprise mock storefront (see Figure 14).

reCAPTCHA Enterprise evaluates a wide range of unusual activities for fraudulent transaction scenarios, including:

- Higher number of basket abandonments
- Lower-than-average basket price
- Higher proportion of failed payment authorizations
- Disproportionate use of the payment step
- Increased chargebacks
- Multiple failed payment authorizations from the same user, IP address, user agent, session, and/or device ID/fingerprint
- Same or similar accounts for both "buyer" and "seller" in sites that facilitate consumer-to-consumer (C2C) commerce
- Increased demand for higher-value goods or services

Because our interaction did not include any known malicious sources and did not deviate in any way from a "normal" transaction scenario, our score was positive (see Figure 15).

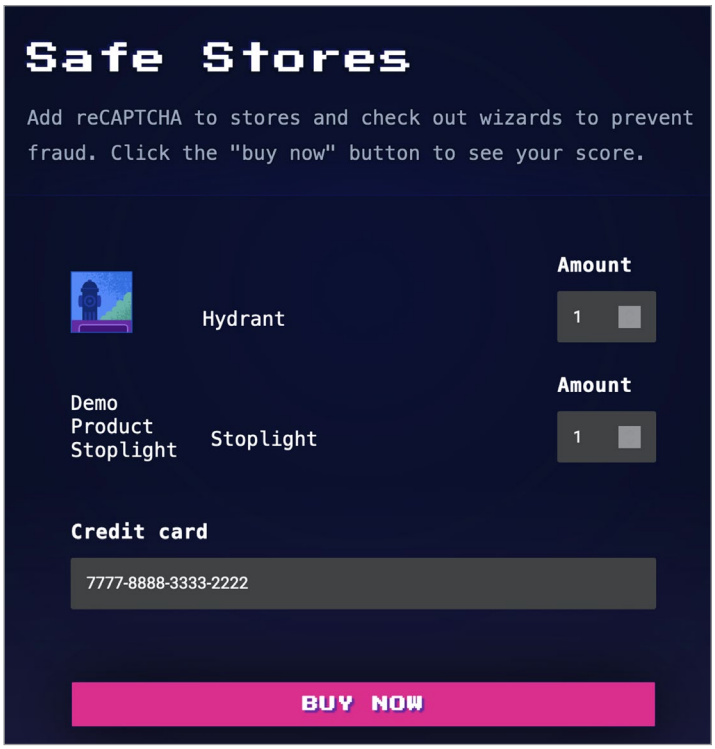

Figure 14. The Sample reCAPTCHA Enterprise Storefront

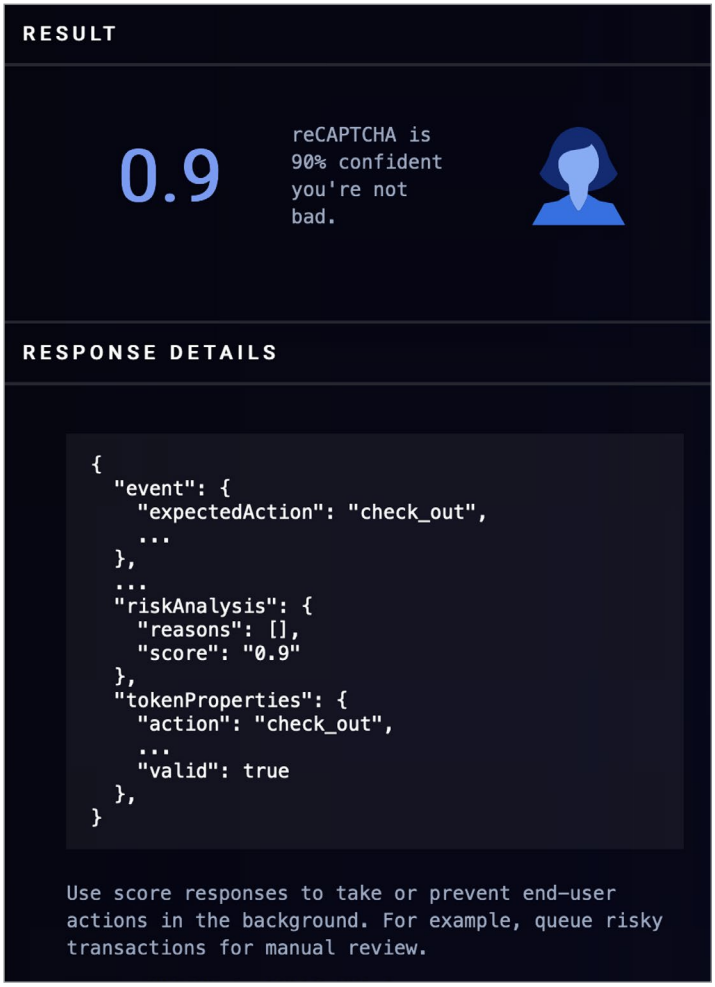

Figure 15. A Normal Payment Checkout Process

In-Depth Fraud Prevention

Deep analytics for fraud prevention is one of the most powerful use cases for reCAPTCHA Enterprise. The **"fraudPreventionAssessment"** component of the platform includes a transaction risk score, as well as multiple fraud scores that identify different types of attacks (see Figure 16).

Let's break down the various elements of this analysis:

- The reCAPTCHA Enterprise Fraud Prevention **"transaction_risk"** score ranges from 0.0 to 1.0, summarizing the risk associated with this transaction from the components below. The risk score 0.0 indicates that the risk is low and the transaction is likely legitimate; 1.0 indicates that the risk is high and the interaction is likely fraudulent.
- The **"cardTestingVerdict"** value detects adversaries using a website to test lists of stolen instruments or brute-force information. When they succeed, this results in potential business loss.
- The **"behavioralTrustVerdict"** value indicates session trust based on behavioral signals on the site and across the internet. This score is particularly helpful if you use an existing fraud-detection engine and want to reduce false positives.
- After sending life-cycle events, including chargeback information, reCAPTCHA Enterprise provides the **"stolenInstrumentVerdict"** value. This detects attacks that may be "low and slow" that are likely to be fraudulent based on the signals that reCAPTCHA Enterprise analyzes on the specific transaction and on the user's behavior across millions of other websites.

We deliberately formulated a suspicious transaction to see the changes in these values, which are shown in Figure 17.

As we observed, the transaction risk was elevated to 1.0, the card testing value was consistent with a stolen payment method, the behavioral trust value was very low, and the "stolen instrument" risk score went up significantly as well.


```
"fraudPreventionAssessment" : {
   "stolenInstrumentVerdict" : {
    "risk" : 0.1
   },
   "transaction_risk" : 0.1,
   "cardTestingVerdict" : {
    "risk" : 0
   },
   "behavioralTrustVerdict" : {
    "trust" : 0.9
   }
}
```


*Figure 16. A Good (Clean) Transaction Analysis by **"fraudPreventionAssessment"***


```
"fraudPreventionAssessment" : {
   "stolenInstrumentVerdict" : {
    "risk" : 0.7
   },
   "transaction_risk" : 1,
   "cardTestingVerdict" : {
    "risk" : 1
   },
   "behavioralTrustVerdict" : {
    "trust" : 0.1
   }
}
```


Figure 17. A Fraudulent Transaction Detected by reCAPTCHA Enterprise

Comment Forms

Comment forms are another common source of online malicious interaction, as attackers can try to post malicious content on sites (cryptocurrency miners, cross-site scripting code, and more) that could lead to user account and system hijacking. We visited the reCAPTCHA Enterprise default comment page (see Figure 18).

As expected, with no malicious input that includes **<script>** tags or other encoded content, our score again remained unchanged (see Figure 19).

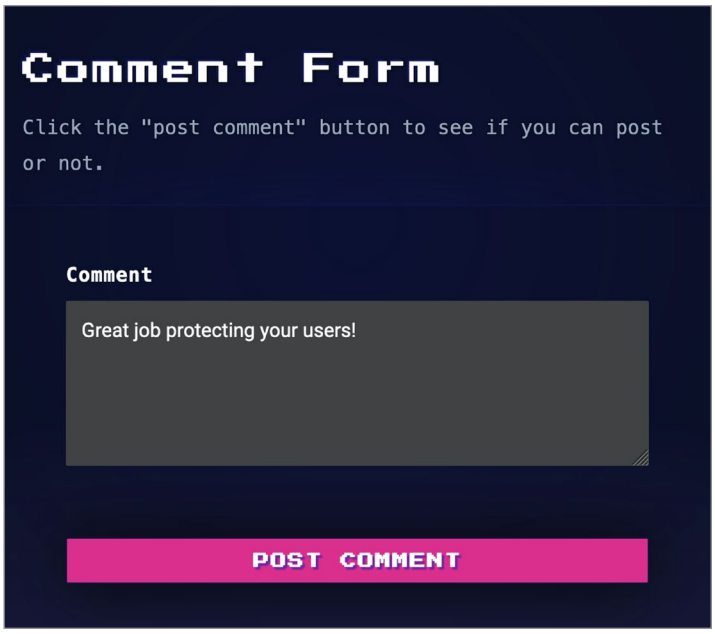

Figure 18. reCAPTCHA Enterprise Comment Page

Figure 19. Nonsuspicious Comment Page Score

Conclusion

As more of our critical services and assets move online, it's important to realize that attackers are shifting their attack strategies to fraudulently interact with sites and services in new, innovative, and perhaps unanticipated ways. For example, one increasingly common tactic called *skewing* focuses on repeated link clicks, page requests, or form submissions to alter metrics and reporting for site interactions and behaviors. To combat this kind of malicious behavior, deep machine learning and threat intelligence are needed in tandem to detect unusual patterns and sources of requests in real time. reCAPTCHA Enterprise offers organizations an entire ecosystem of tools for both detecting and responding to fraud. These tools leverage Google's extensive experiences and observations to help reverse the current attack trends. Increasingly, automation and integration features of reCAPTCHA Enterprise, such as flexible APIs, client libraries, and application integration components like response tokens, are becoming more of a must-have solution for more cloud-hosted applications and services than ever before, especially because they are completely invisible and frictionless for the legitimate end users of these applications.

Sponsor

SANS would like to thank this paper's sponsor:

Google Cloud