ForgeRock®

White Paper

Block the Bots to Keep Your Customers Coming Back

Overview

Retail purchasing bots pretending to be customers buy up supplies of sought-after products and resell them at a markup. This results in your loyal customers missing out on products they are eager to buy. This is happening with increasing frequency. In this white paper, we look at an example of a popular retail bot named Birdbot to understand how it works and to learn about more destructive bots so you can learn how the ForgeRock Identity Platform defends against bots today and in the future, and to help ensure that your loyal customers have a great experience.

The Problem

As shopping has increasingly moved online over the last decade, there has been a seismic shift in how we engage with retailers. And, there's been a corresponding impact on one of the oldest rules of retail: that supply and demand have direct consequences on pricing.

Let's look at this rule in a real-life example. Before the days of eCommerce, many people lined up to be the first at the box office to purchase tickets for their favorite musician. Sometimes this meant lining up for hours or through the night. The most dedicated fans could usually secure concert tickets at the price set by the initial ticket seller. Today, concert ticket sales typically begin online at a specific time. Fans desperate for tickets are often on the website before the sale time, refreshing the web page constantly as the sale time approaches. The time arrives, eager fans hope to buy tickets, and the site crashes. While a lucky few may manage to check out successfully, most fans won't. Within minutes, most tickets will have sold out completely.

What happened? Most likely two things have occurred:

1. The site couldn't cope with the sudden huge spike in traffic.

2. Bots bought all the tickets to resell on an online marketplace and other online sites.

Variants of this scenario happen frequently, such as with new movie releases, limited edition products, or the hottest video game. During the COVID 19 pandemic, automated bots swept up many of the online supplies of the Nintendo Switch to resell on an online marketplace at significant markup. The problem is becoming so significant that websites and services are popping up to alert customers when popular or hard-to-find products are in stock.

This is happening everywhere, all of the time. It's impacting the customer experience and leaving retailers to deal with frustrated customers as well as processing numerous returns if the scalpers fail to resell the items and make a profit. Only the scalpers benefit from this trend.

Retailers need to be able to determine if their products are being purchased by a legitimate, loyal, excited customer that has been eagerly waiting for a new product release or by a cold, soulless bot from the darknet.

Copyright © 2020 ForgeRock, All Rights Reserved.

Solution

The ForgeRock Identity Platform can help you identify and stop bots without accidentally driving your real customers away. Before we dive into how we do that, we want to ensure that your site can handle the sudden spike in traffic that a product launch or similar event will generate. This typically means leveraging additional infrastructure. To avoid paying for extra capacity after traffic has subsided back to normal levels, you can employ cloud-based elastic scalability options. The ForgeRock Identity Platform has been built for high scale, and, with modern DevOps tools like Kubernetes and Docker, it can be automatically deployed to scale according to demand. You can find more information here.

People Versus Bots: How Do Bots Work?

It's critical that you determine if a site visitor is a real person or a purchasing bot. To do this, you need to examine as much contextual information as possible and use it to make decisions. The key is to avoid making the process onerous for your real customers, or you risk them not coming back next time. This involves striking a careful balance between customer experience and the required security friction to keep your site safe.

Deckard: "She's a replicant, isn't she?"
Tyrell: "I'm impressed. How many questions does it usually take to spot them?"

- Blade Runner

Introducing Birdbot

It's useful to understand how automated bots work. One of the most prolific retail bots active today is called Birdbot. It's freely available and even well documented. It was designed to automate purchasing and checkout at Walmart and Best Buy in the U.S. Birdbot was used to purchase Nintendo Switch consoles for online resale during the early stages of the pandemic.

How does it work?

1. Users download Birdbot and run it on their local computers.
2. Users configure Birdbot with:
   a. *Tasks:* Items they want to purchase.
   b. *Profiles:* Shipping, billing, and payment information.
   c. *Proxies:* List of internet proxies that Birdbot should use. A proxy is a server that allows users to route their internet traffic to hide its real origin. There are legitimate uses for proxies - such as ensuring user privacy, particularly in regions where harsh internet usage laws are in place - but in this case, they are used to avoid detection. By using proxies, Birdbot hides the user's origin. This prevents the retailer's website from easily blocking it, as different internet addresses are used all the time.

> The proxies page is where you set up proxies. The bot uses these proxies to avoid ip bans and bot protection blocking.

Figure 1: Excerpt from Birdbot documentation

   d. *Settings:* Primarily where Birdbot should send notifications of success or failure.

Copyright © 2020 ForgeRock, All Rights Reserved.

3. Birdbot mimics a normal web browser by configuring a user-agent. User-agents are how browsers, phones, and other devices identify themselves so a website knows how to format the content to best fit that device. Mimicking the user-agent isn't difficult, as they aren't designed with security in mind.

```
def monitor(self):
    headers = {
        "accept": "text/html,application/xhtml+xml,application/xml;q=0.9,image/webp,image/apng,*/*;q=0.8,application/signed-exchange;v=b3;q=0.9",
        "accept-encoding": "gzip, deflate, br",
        "accept-language": "en-US,en;q=0.9,zh-CN;q=0.8,zh;q=0.7",
        "cache-control": "max-age=0",
        "upgrade-insecure-requests": "1",
        "user-agent": "Mozilla/5.0 (Macintosh; Intel Mac OS X 10_15_1) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/81.0.4044.69 Safari/537.36"
    }
```

Figure 2: Excerpt of Birdbot Code for User-Agent

4. Birdbot then runs code that allows it to mimic user interactions, including checking out, submitting payment information, and choosing shipping options.

```
def submit_shipping(self):
    headers={
        "accept": "application/com.bestbuy.order+json",
        "accept-encoding": "gzip, deflate, br",
        "accept-language": "en-US,en;q=0.9,zh-CN;q=0.8,zh;q=0.7",
        "content-type": "application/json",
        "origin": "https://www.bestbuy.com",
        "referer": "https://www.bestbuy.com/checkout/r/fulfillment",
        "user-agent": "Mozilla/5.0 (Macintosh; Intel Mac OS X 10_15_1) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/81.0.4044.92 Safari/537.36",
        "x-user-interface": "DotCom-Optimized"
```

Figure 3: Excerpt of Birdbot Code for Submitting Shipping Method

5. Birdbot does not log in. Instead, it utilizes guest purchasing.

6. Birdbot repeats this process many times a minute. It checks if something is in stock and automatically purchases it when it is.

Birdbot is easy to use, well built, and gets the job done. But there are even more sophisticated, disruptive, and damaging bots in use today. As retailers push back against simple purchasing bots, these more sophisticated bots will continue to be developed and become increasingly more powerful. The game of Whac-A-Mole will continue.

Copyright © 2020 ForgeRock, All Rights Reserved.

Attacks from Botnets

The bots that cause the disruptive events that reach the news are typically part of a bot network, or "botnet." Botnets are groups of thousands to hundreds of thousands of devices that have been infected or compromised by malware. The malware hides on the devices and is controlled through a central command server that directs the bots to perform whatever actions the attacker requires of them. The advantage of a botnet is that it is distributed. Each bot is hidden inside a unique device connected to the internet, and these devices are distributed all over the world. This network of bots is effectively an evolution of the proxy idea utilized by simpler bots. Proxies usually only offer a range of internet addresses to hide behind, whereas a botnet can span the entire internet.

With this level of connection and coverage, the botnet can be directed to overwhelm a target – like a retailer's website – with more requests than it can cope with. In what is typically known as a DDoS (Distributed Denial-of-Service) attack, the website is unable to serve legitimate requests and all traffic is effectively frozen. The reason website owners are so vulnerable to DDoS attacks is that the attacking requests originate from many different internet addresses, so it's not sufficient to simply block traffic from a specific location.

Unfortunately, retailers are also vulnerable to attacks beyond bot-enabled reselling and DDoS attacks. Some bots put items into a basket without checking out, denying them to legitimate shoppers as a form of malicious attack. Others engage in mass credential "stuffing," using breached password databases to make thousands of logon attempts until they find a username and password combination that works.

If left unchecked, bots – whether simple or sophisticated – can have a significant impact.

Fighting Back

So what can retailers do? To recap, we know that bots like Birdbot:

- Typically use proxies or a botnet to hide their real origin.
- Pretend to be a normal browser and interface with the website using code to mimic human interaction.
- Do not tend to login, but instead make use of guest checkout features.

Of course, not all bots will behave in the same way, and we can fully expect bots to evolve as retailers adapt to combat them. In the meantime, though, this is a good set of assumptions for us to start with. Let's look at how the ForgeRock Identity Platform empowers you to defend against these tactics today and into the future.

Copyright © 2020 ForgeRock, All Rights Reserved.

What Does Bot Traffic Look Like Compared to Regular User Traffic?

Let's first examine what typically happens when you shop online. For simplicity, we'll assume users are shopping directly on a retailer's website rather than on a native mobile app.

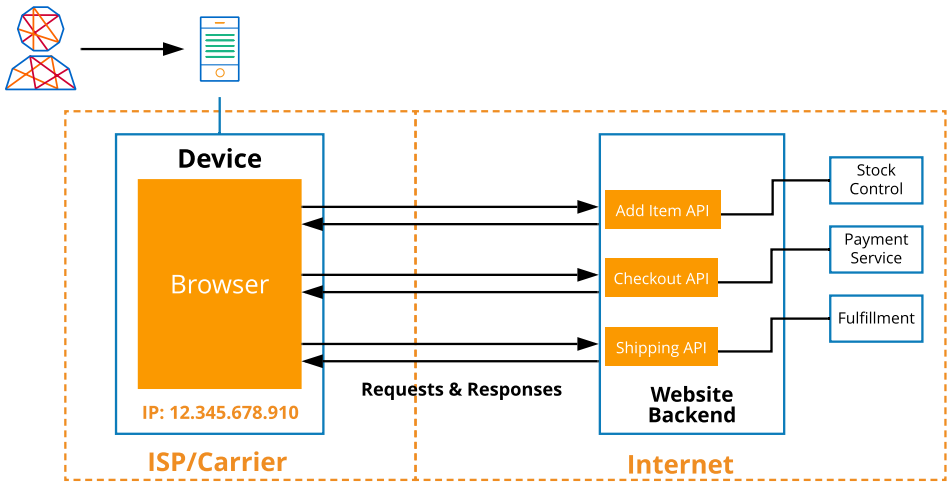

Figure 4: Visualization of typical user purchase process

When shopping online, users typically:

1. Access their device – a mobile phone, tablet, or computer – that has already been issued an IP address by the user's Internet Service Provider or mobile carrier.
2. Open their browser of choice and enter the URL of the website they want to visit.
3. Browse around, search for items to buy, and put them into their shopping cart.
4. Enter shipping and payment information to check out and finalize their purchase.

What's happening under the hood is a set of internet requests and responses between the device and the retail website. As users put things into the shopping cart and go through the check out process, the user interface will invoke application programming interfaces (APIs) on the website backend that in turn call out to other systems to make it all happen.

> IP Addresses: Internet Protocol (IP) addresses are similar to postal addresses that enable you to send and receive mail. IP addresses enable devices to send and receive internet traffic. They are typically issued by Internet Service Providers (ISPs) or mobile carriers and look similar to the following: 212.58.244.57.

Copyright © 2020 ForgeRock, All Rights Reserved.

How does this flow change when a bot like Birdbot is being used?

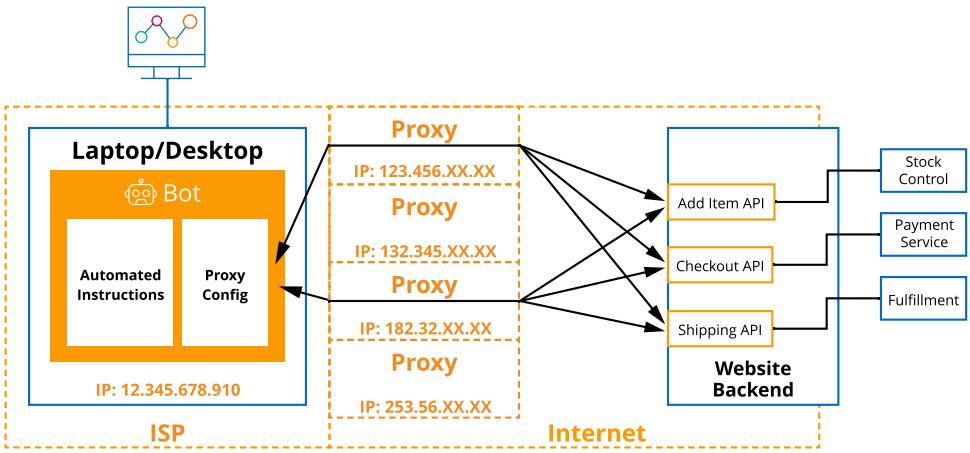

Figure 5: Visualization of bot-driven purchase process

There are a few key differences:

- The bot user has most likely set the bot to run and then left it running. They are no longer interacting with your website directly.
- The bot is not using the regular user interface that the web browser provides. Instead, the bot is invoking APIs directly.
- The bot is masking its origin using proxies.

Protecting Access with Gateway and Identity Context

To combat bot traffic, you need to identify whether traffic is likely to be originating from a bot or a device being used by real human customers so you can decide how to handle and authenticate it. ForgeRock Identity Gateway and Intelligent Access solutions help you do just that.

A gateway examines internet traffic trying to access your website and decides whether or not to let it through. Its decision is based upon whether the traffic has been issued a token, commonly referred to as a session. Gateways examine sessions and either allow or deny traffic to your site based on the determination of session validity. They can also protect APIs. The ForgeRock Identity Platform includes Identity Gateway and the platform can also be integrated with other commercial gateways like those from Apigee, Mulesoft, and Kong.

Copyright © 2020 ForgeRock, All Rights Reserved.

Once you have the gateway in place, you need to ensure that your website is only issuing sessions to real customers and not bot-generated traffic. This is where ForgeRock Intelligent Access comes into play. Intelligent Access enables you to design user journeys consisting of different steps or nodes. These collect data and make decisions based upon that data. When the user journey is completed, a session is issued if Intelligent Access determines the traffic is from a legitimate customer. Sessions are signed cryptographically to ensure they cannot be tampered with.

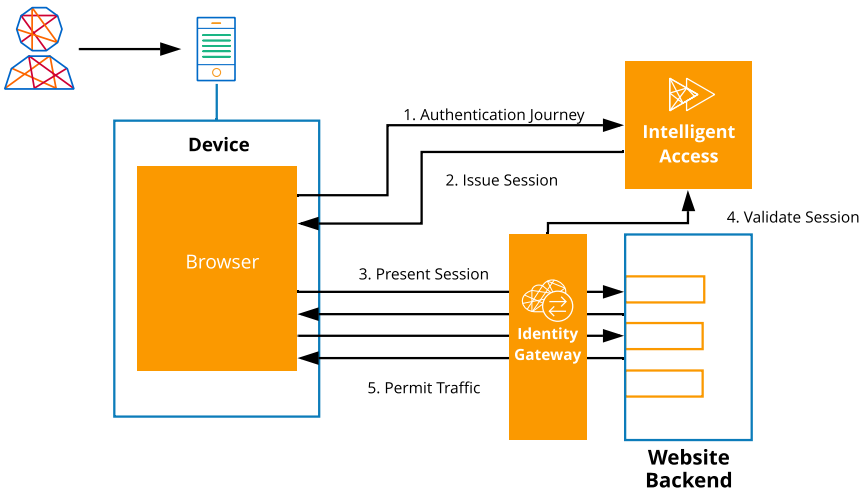

Figure 6: Using Intelligent Access and Identity Gateway

Traditionally, sessions were automatically issued when a user logged in after entering a username and password, but this simple logical workflow is no longer sufficient. Usernames and passwords have since been augmented with multi-factor authentication and biometrics. And, with features like guest checkout, it's no longer as simple as your customers logging in or out of a site. ForgeRock Intelligent Access is flexible, enabling you to define custom journeys that take these different situations into account.

Creating Secure User Journeys

With ForgeRock, you can define a new, secure user journey that lets real customers in while keeping bots out. You can do this by leveraging comprehensive out-of-the-box capabilities and, where required, integration with ForgeRock's Trust Network Technology Partners. When the user journey is complete and Intelligent Access determines the visitor is legitimate, it will issue a session. The session can then be presented to the Identity Gateway alongside the internet traffic from the device. The session is checked for validity and only then is the internet traffic allowed to pass through to the website or API.

In the ForgeRock Identity Platform, user journeys are implemented using building blocks called nodes. Each node fulfills a simple function, and using the intuitive and visual user interface (UI) it's simple to connect them together to form different user journeys.

Copyright © 2020 ForgeRock, All Rights Reserved.

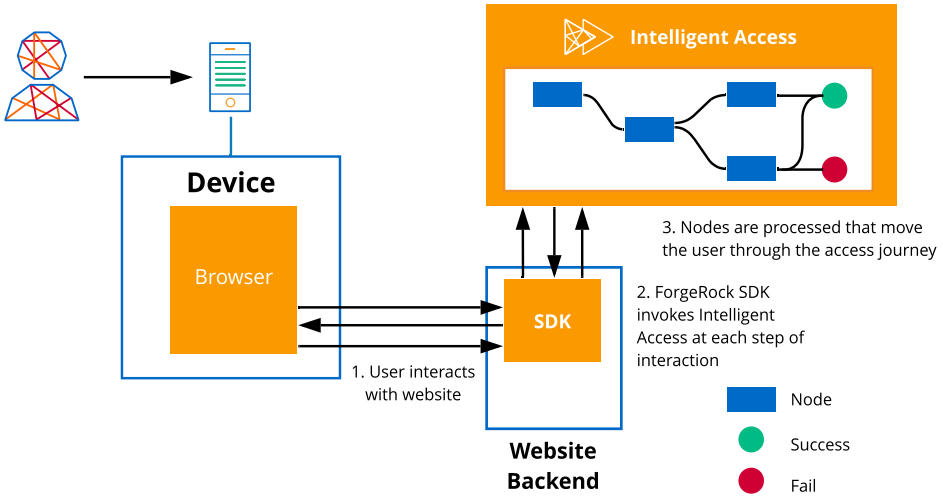

Figure 7: The customer journey and Intelligent Access

Using nodes enables you to define customer journeys that take multiple factors into consideration, including:

IP Addresses: Because bots hide their origin, you cannot simply block a particular internet address. They also modify the user-agent and manipulate other identifying information relating to their connection. There are many third-party security services that actively monitor the internet for bot activity and collate lists of the IP addresses where such bots are known to originate. ForgeRock Intelligent Access can drop a node into the user journey that calls out to these services in real time and makes a decision based on the result.

reCAPTCHA: To confirm that you're dealing with a bot instead of a real human customer, you can request the visitor completes a reCAPTCHA. Bots tend to struggle with these. However, reCAPTCHA can be frustrating for legitimate customers, especially if they are trying to check out quickly. It is a good practice to only deploy reCAPTCHA when there is a high likelihood you are encountering a bot and not as a blanket practice for all site visitors (which we explain below).

Device and Environmental Context: Real users use devices, and those devices provide crucial contextual information. We know that Birdbot and other bots modify user-agents to mimic browsers. With user consent you can scan the device – such as laptops, mobile phones, computers or tablets – and gather information such as user-agent and other information to build a device fingerprint. A simple pop-up message can request the user's consent to share information such as location. If the user consents, this gives you another data point and increases your assurance that the user is real and not a bot.

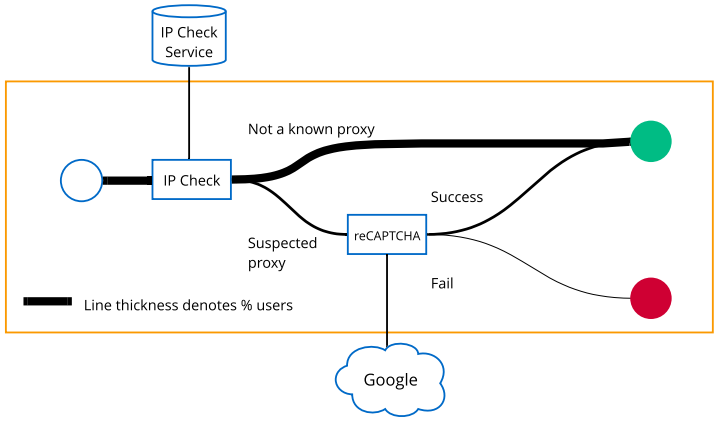

Figure 8: Intelligent Access journey which only enforces reCAPTCHA when appropriate

Copyright © 2020 ForgeRock, All Rights Reserved.

Logging In: Encouraging your customers to log in enables them to build up a profile that allows you to match their devices to their identity. We certainly don't want to force customers to log in – many people enjoy the freedom, privacy or speed of guest checkout – but there are ways to encourage users to log in. You can use social media sign-ins through Facebook or Google. And, you can employ ForgeRock Go, a solution that does away with usernames and passwords entirely while still allowing you to build profiles and increase confidence in the legitimacy of your users.

Behavioral Biometrics: Machine learning and artificial intelligence (AI) have opened up new approaches to determining if you're dealing with a real human, a real human you've interacted with before, or a bot. ForgeRock Trust Network Technology Partners such as Biocatch and SecuredTouch, available in the ForgeRock Marketplace, examine how the user interacts with the web page and determine if the user is the same user who has used that account previously, a real person that you recognize, or a bot. These services can offer a high level of assurance and can analyze behavior in a way that is completely transparent to the user and does not impact their experience.

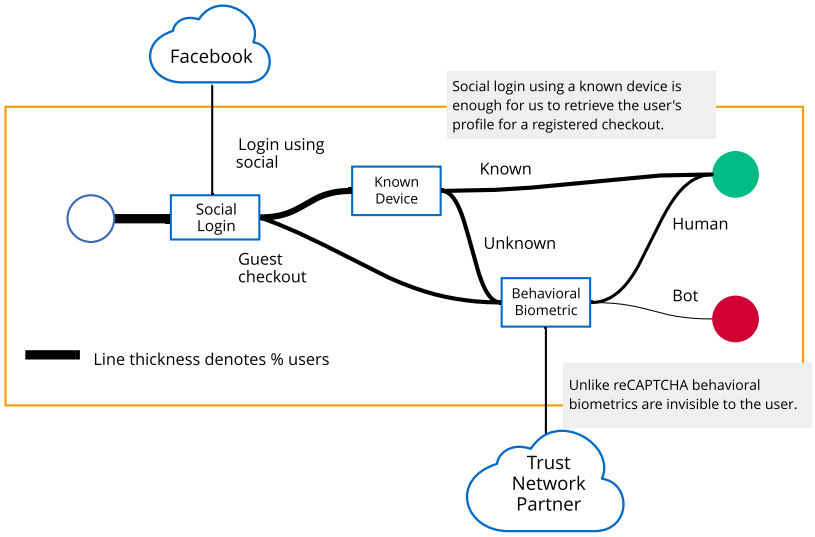

Figure 9: Intelligent Access journey which minimises friction

T-800: "I now know why you cry, but it is something that I can never do."

- Terminator 2: Judgement Day

Customer Activity, Shopping Seasons, and Product Releases: By enabling the use of scripting throughout the nodes that comprise the user journey, ForgeRock Intelligent Access lets you incorporate additional valuable information. For example, you can create a node that consults a list of upcoming release dates, stock refreshes, or peak shopping seasons to turn on and off measures as appropriate. You can also look at the number of orders made by a particular device fingerprint and enforce a soft limit or further validation through reCAPTCHA. This flexibility and intelligence further increases your contextual information and assurance.

Metrics and A/B Testing: The ForgeRock Identity Platform enables you to easily conduct A/B testing and ensure that you're classifying legitimate and bot traffic with accuracy and not accidentally driving real customers away. For example, you can apply new journeys to a small subset of test customers to gauge results and make sure it doesn't increase cart abandonment. Metrics are built into nodes so you can instrument the journey along the way and understand where the journey can be tweaked and improved.

ForgeRock recognises that no organisations are exactly alike and the ForgeRock Identity Platform gives you the flexibility to define user journeys that exactly fit your requirements.

Copyright © 2020 ForgeRock, All Rights Reserved.

Bringing It All Together

While there is no silver bullet to defeating all retail bots, ForgeRock and our partners can help you deploy smart, future-proofed approaches that mitigate a majority of the problems they can cause. Intelligent Access helps you build user journeys that quickly identify real users and channels them to your site without delay while simultaneously blocking bots. If you prefer, you can also redirect bots to honeypots so you can learn their secrets in a safe, sandboxed environment.

The Future

Bot Adaptation

Security is like an arms race. As you learn to respond to threats, bad actors adapt exploits – and you need to adjust your response. You need a solution that allows you to keep up with bots as they adapt. In addition, authentication is no longer a simple, single, one-time event with a binary result. Instead, it's a nuanced spectrum of different conditions that requires continuous review and assessment.

Worf: "They've adapted."

- Star Trek: First Contact

"All of this has happened before, and all of this will happen again."

- Battlestar Galactica

ForgeRock Intelligent Access incorporates the concepts of Gartner Group's Continuous Adaptive Risk and Trust Assessment (CARTA) framework. The ForgeRock Identity Platform allows you to react and adapt as bots do, leveraging best-of-breed solutions from our Trust Network partners to rapidly counter new threats. With it, you can drop new nodes into your user journey, wire them up in the simple drag-and-drop user interface, and stop the bots. Besides saving many hours of development time and resources, you gain confidence, knowing you are relying on proven solutions.

Conclusion

Bots are here to stay. And you need to stay ahead of them to ensure your real customers are having a safe and positive retail experience with your brand. The ForgeRock Identity Platform empowers you to deliver that experience today, tomorrow, and into the future. Contact us to learn more.

About ForgeRock

ForgeRock, the leader in digital identity, delivers modern and comprehensive Identity and Access Management solutions for consumers, employees and things to simply and safely access the connected world. Using ForgeRock, more than a thousand global customer organizations orchestrate, manage, and secure the complete lifecycle of identities from dynamic access controls, governance, APIs, and storing authoritative data – consumable in any cloud or hybrid environment. The company is privately held, and headquartered in San Francisco, California, with offices around the world. For more information and free downloads, visit **www.forgerock.com** or follow ForgeRock on social media.

Follow Us

ForgeRock

Copyright © 2020 ForgeRock, All Rights Reserved.