onelogin

WHITEPAPER

The Ultimate Guide to Risk-based Authentication

Passwords are not enough

As you probably know, protecting applications with just a password is a recipe for breaches. What you might not know is that the situation is getting worse. The 2016 Verizon Data Breach Investigations Report (DBIR) stated that nearly two-thirds of data breaches involve weak, default, or stolen passwords; the 2017 edition of the DBIR showed that number increasing to over 80%. Two factors that may be driving this trend: weak passwords are common, and those weak passwords can be guessed in less than a millisecond.

Static rules are not enough, either

Because of the weakness of passwords, multi-factor authentication has become critical for strengthening security. But traditional access managers use static rules that can't keep up with today's constantly evolving security risks. These static rules lead them to not prompt for MFA during logins that are high risk. And conversely, to require an additional factor for low-risk logins. This leads to the worst of both worlds: increased friction for users, and increased risk for organizations.

What is required is to replace static rules with machine intelligence that can better measure the risk of a login and prompt for MFA when necessary. We call this Adaptive Authentication. In this whitepaper, we describe use cases where static rules fail to measure login risk accurately, and thus fail to deploy MFA when truly needed. We also describe how Adaptive Authentication better measures login risk. We conclude with how to move forward with Adaptive Authentication in your organization.

How static rules fail

Here are three examples of how static rules fail to measure login risk accurately.

Failure 1: Phishing attacks

Phishing—sending a fraudulent email enticing the recipient to click a link or attachment that typically installs malware or capture

Content

credentials—is an incredibly common way for hackers to attack organizations. Over 90% of breaches in the 2017 Verizon DBIR involve some sort of phishing. Unfortunately, static rules are poorly equipped to address this kind of threat. This is because static rules often involve IP whitelisting, which treats all company IP addresses as trusted and thus not requiring MFA. The thinking goes, If login is coming from company premises, it is not high risk.

But consider a case where a phishing email contains a link to install malware on an employee laptop, and this malware then attempts to access SaaS applications using a brute force attack consisting of repeatedly trying many different passwords. Let's assume that some of those applications haven't been hardened to deal with these types of attacks via some login throttling, account lockouts, captchas, or other measures. Eventually, the malware finds the right password. When it does, static rules will not challenge for MFA due to IP whitelisting. The malware accesses the SaaS applications, then can continue its attack.

Over 90% of breaches in the 2017 Verizon DBIR involve some sort of phishing. But, static rules are poorly equipped to address this kind of threat.

Failure 2: Malicious visitors

Look around your office. Is everyone onsite a trusted employee? Of course not. In your office, you have contractors, customers, vendors, job candidates, and service staff. Any one of them could discover your office WiFi password with a bit of social engineering: just pretend the guest WiFi isn't accessible, you urgently need network access, and a well-meaning employee might provide it to help out.

When companies use static rules to setup IP whitelisting, however, the results can be disastrous. Our malicious visitor can then attempt to login to company applications without providing a second authentication factor using just a (weak) password.

Failure 3: Trusted remote employees

Consider an employee that has been working from their home office for months. Every day, they have the same IP address, physical location, laptop, and work hours. Also, they consistently provide a second factor to authenticate. Authentication products that rely on static rules will always consider this a high-risk login and force that employee to provide a second factor before logging on. The authentication system won't learn after a number of successful authentications that this is a low-risk situation.

This is an example of "security theater," a practice that might seem to increase security but actually doesn't. It hurts employee productivity by requiring them to repeatedly provide a second factor. It actually reduces security by adding superfluous security events to your Security Events and Information Management (SIEM) system. This adds noise that makes it harder to detect actual threats, which matters because an enterprise might see 200,000 security events a day. Adding to this firehose of data

increases both the time for a security analyst to investigate security incidents and the time for a hacker to execute an attack. Garbage in, Garbage out.

Machine learning for smarter risk scoring

Because static rules fail to measure risk accurately, we need a better way to determine when MFA should be required, and when it shouldn't. That's one of the things we've done with OneLogin Adaptive Authentication. Built for security teams who need to reduce the risk of account takeovers, Adaptive Authentication uses machine learning to score the risk of each login attempt and challenges users making high-risk logins to provide an additional authentication factor. These can include a one-time password (OTP) or security questions.

OneLogin Adaptive Authentication is risk-based MFA. It uses a broad set of inputs—including networks, geography, devices, and time—to learn a user's typical usage patterns or behavior over time. The greater the deviation from their typical usage, the higher the risk score. Risk scores will also be increased by network traffic coming in from an IP address with low trust (such as Tor exit relay), a new country or city, or from locations that are too far apart to be accessed by a single person within a given timeframe.

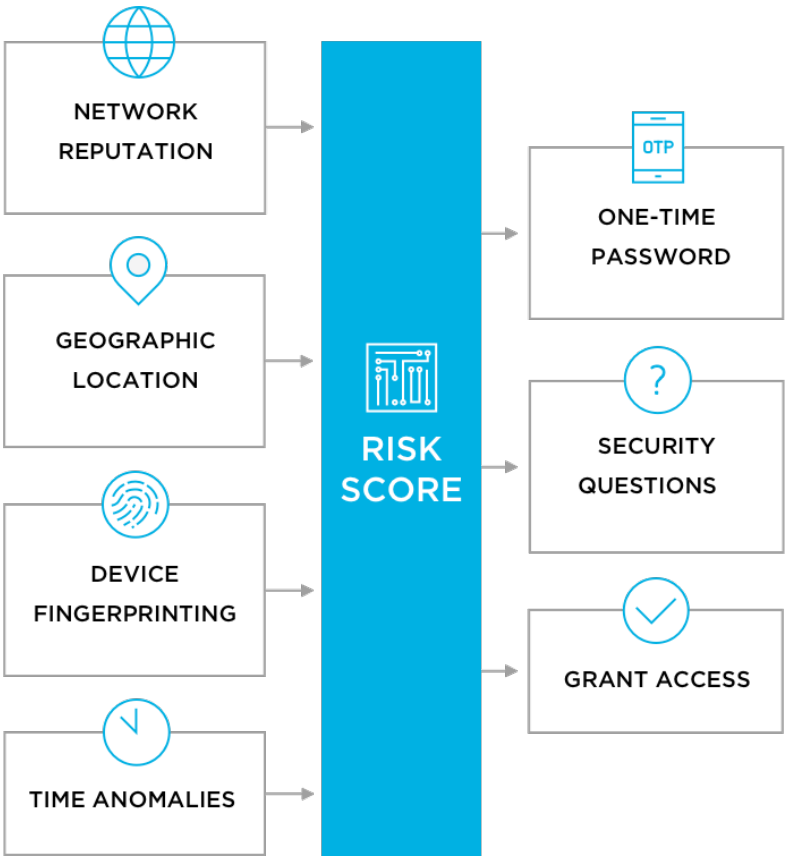

Jargon Buster:

What is Tor?

Tor can be thought of as a randomized virtual private network. When you connect your computer to Tor, it routes your packets to appear on the Internet in a random location in the world. Your connection to the Internet might be halfway around the world, or the next town over. Where your connection appears on the Internet is called a Tor exit relay.

Tor enables users to anonymously access the Internet. Unfortunately, that same anonymity is useful to hackers looking to mask their identity during attacks. For this reason, OneLogin Adaptive Authentication treats logins from Tor exit relays as high risk.

What is AlienVault Open Threat Exchange?

Open Threat Exchange (OTX) is a crowdsource computer security platform, and is run by AlienVault. OTX enables security professionals to work together to share information about viruses, malware, and other cyber attacks. Networks listed as compromised by OTX are treated as untrusted by OneLogin Adaptive Authentication.

What is Project Honeypot?

Project Honeypot is a distributed system to identify spammers and their spambots that scrape email addresses off websites. It collects IP addresses of known spambots, which can then be treated as untrusted networks by OneLogin Adaptive Authentication. .

Unlike static rules, OneLogin Adaptive Authentication gets smarter over time, enabling it to accurately score risks in a variety of scenarios, including:

- Phishing attacks
- Malicious visitors
- Trusted remote employees
- Mobile device thieves
- Anonymized hackers
- Malicious botnets
- Account takeovers

How Adaptive Authentication Addresses Threats

To demonstrate how Adaptive Authentication works check out these seven situations below:

Scenario 1: Phishing attacks, revisited

Let's revisit our previous example of a phishing attack that installs malware on an employee computer, running on a company network, to repeatedly attempt to login to company apps. For simplicity, let's assume that the malware only attempts logins during office hours to avoid detection, and that there is not an employee logged in with the same account the malware is attacking. After many login attempts, when the malware finally uses the correct password, here's how Adaptive Authentication would view the risk profile of this authentication attempt:

Risk Factor	Result	Risk Level
Network Reputation		
Flagged in AlienVault Open Threat Exchange	✖	↓
Flagged in Project HoneyPot	✖	↓
Access from Tor Network	✖	↓
New IP address	✖	↓
Geographic Location		
New city or country	✖	↓
Device Fingerprinting		
New device	✔	↑
New or infrequently used OS	✔	↑
New browser	✔	↑
Time Anomalies		
Unusual time of day	✖	↓
Fast user movement	✖	↓
Simultaneous access from two distant locations	✖	↓

Legend: ↑ High Risk ↓ Low Risk

Because this login is coming from the company network, it does not trigger any risks associated with coming from a disreputable network or geography. But the malware's HTTP client is a new "device" that we haven't seen before, it has a new device fingerprint, which indicates high risk. So, Adaptive Authentication challenges the malware to authenticate using two-factor authentication. The malware, being software, obviously can't exactly pull out a phone and submit a one-time password. For this reason, the login is blocked, and the phishing attack is stopped.

This brings up an important point about Adaptive Authentication: it doesn't have to trust a login attempt just because it's coming from a company IP address, inside a firewall. It holds login attempts to a higher standard: they need to be scrutinized for their device fingerprint and any time anomalies. Because of this, you can think of Adaptive Authentication as implementing the concept of Zero Trust for identity.

Scenario 2: Malicious visitors, revisited

Back to our malicious visitor example, where a hacker gets your office WiFi password, then attempts to use an employee's credentials to access an app during office hours. Let's assume the visitor is using a laptop that Adaptive Authentication has not previously seen, and that they are using an account that an employee is not currently logged into.

Adaptive Authentication would view the risk profile of this authentication attempt similarly to the previous phishing example:

Because the malicious visitor's laptop has never been seen before, it has a new device fingerprint, which indicates high risk and prompts for MFA, blocking the visitor's attempts to access employee accounts.

Risk Factor	Result	Risk Level
Network Reputation		
Flagged in AlienVault Open Threat Exchange	✖	↓
Flagged in Project HoneyPot	✖	↓
Access from Tor Network	✖	↓
New IP address	✖	↓
Geographic Location		
New city or country	✖	↓
Device Fingerprinting		
New device	✔	↑
New or infrequently used OS	✔	↑
New browser	✔	↑
Time Anomalies		
Unusual time of day	✖	↓
Fast user movement	✖	↓
Simultaneous access from two distant locations	✖	↓

Scenario 3: Trusted remote employees, revisited

Risk Factor	Result	Risk Level
Network Reputation		
Flagged in AlienVault Open Threat Exchange	✖	↓
Flagged in Project HoneyPot	✖	↓
Access from Tor Network	✖	↓
New IP address	✖	↓
Geographic Location		
New city or country	✖	↓
Device Fingerprinting		
New device	✖	↓
New or infrequently used OS	✖	↓
New browser	✖	↓
Time Anomalies		
Unusual time of day	✖	↓
Fast user movement	✖	↓
Simultaneous access from two distant locations	✖	↓

Let's now revisit the example of a longtime home office worker who has repeatedly logged in from the same IP address, physical location, and laptop, during the same set of hours, and has consistently provided a second factor to authenticate. Once this pattern has been established, here's how Adaptive Authentication would view the risk profile of this authentication attempt:

Because this login follows a well-established usage pattern, there are low risk levels related to network reputation, geographic location, device fingerprint, and time anomalies. The home office worker could be allowed to log in with just their account ID and password. Additionally, SIEM systems, such as Splunk, Sumo Logic, and ELK, wouldn't have the noise of false positives from remote employees logging in as usual.

Static whitelisting is unworkable for trusted remote employees. If there are many of these remote employees, whitelist rules will grow complex and unwieldy to maintain, especially if they change residences. Worse, this complexity can lead to security gaps if employees depart and their whitelisted IPs are not removed from the static rules.

Scenario 4: Mobile device theft

Suppose an employee's unlocked mobile device is stolen one evening after business hours, and shortly afterward, the thief attempts to use it to log into corporate apps. Here's how Adaptive Authentication would view the risk profile of this authentication attempt:

Because this device has been used to access company apps in the past, there is a low risk associated with the device fingerprint. However, the thief is using the device during the middle of the night, an unusual time that drives up the risk score. The thief is also coming in from a previously unseen network and city; these further increase the risk score. Given the high risk, the thief is challenged for MFA and is not able to access company apps.

Risk Factor	Result	Risk Level
Network Reputation		
Flagged in AlienVault Open Threat Exchange	✖	↓
Flagged in Project HoneyPot	✖	↓
Access from Tor Network	✖	↓
New IP address	✔	↑
Geographic Location		
New city or country	✔	↑
Device Fingerprinting		
New device	✖	↓
New or infrequently used OS	✖	↓
New browser	✖	↓
Time Anomalies		
Unusual time of day	✔	↑
Fast user movement	✖	↓
Simultaneous access from two distant locations	✖	↓

Scenario 5: Anonymized hackers

Risk Factor	Result	Risk Level
Network Reputation		
Flagged in AlienVault Open Threat Exchange	✖	↓
Flagged in Project HoneyPot	✖	↓
Access from Tor Network	✔	↑
New IP address	✔	↑
Geographic Location		
New city or country	✔	↑
Device Fingerprinting		
New device	✔	↑
New or infrequently used OS	✖	↓
New browser	✖	↓
Time Anomalies		
Unusual time of day	✖	↓
Fast user movement	✖	↓
Simultaneous access from two distant locations	✖	↓

Suppose a hacker was attempting to access corporate apps from a remote location, during business hours, using their own device. To remain anonymous, they protect their identity by using Tor. Let's assume the hackers have correctly guessed the OS and browser typically used by this user. Adaptive Authentication would view the risk profile of this authentication attempt as:

Coming from a Tor exit relay would drive risk score higher, as would coming from a previously unseen IP address, city, and device. This hacker would be challenged to authenticate using multi factor authentication, blocking them from company apps.

Scenario 6: Malicious botnets

Suppose a hacker uses malicious software running on multiple nodes on a botnet to attempt to perform a brute force password attack on company apps, attempting to log in during business hours. As with the previous example, let's assume the hackers correctly guessed the user's typical OS and browser. Adaptive Authentication would view the risk profile of these repeated authentication attempts as:

Risk Factor	Result	Risk Level
Network Reputation		
Flagged in AlienVault Open Threat Exchange	✔	↑
Flagged in Project HoneyPot	✔	↑
Access from Tor Network	✖	↓
New IP address	✔	↑
Geographic Location		
New city or country	✔	↑
Device Fingerprinting		
New device	✔	↑
New or infrequently used OS	✖	↓
New browser	✖	↓
Time Anomalies		
Unusual time of day	✖	↓
Fast user movement	✔	↑
Simultaneous access from two distant locations	✖	↓

Adaptive Authentication would detect a number of issues. First, the IP would likely be listed as malicious by AlienVault Open Threat Exchange and/or Project HoneyPot. Additionally, the IP address, city, device, and possibly country would be one that we haven't seen previously, all of which would increase risk score. Finally, if two geographically remote botnet nodes tried to login, Adaptive Authentication would detect that as impossibly fast user movement, further driving risk higher. All of these would trigger the use of MFA, which would block access by the malicious software.

Scenario 7: Account takeovers

Risk Factor	Result	Risk Level
Network Reputation		
Flagged in AlienVault Open Threat Exchange	✖	↓
Flagged in Project HoneyPot	✖	↓
Access from Tor Network	✖	↓
New IP address	✔	↑
Geographic Location		
New city or country	✔	↑
Device Fingerprinting		
New device	✔	↑
New or infrequently used OS	✔	↑
New browser	✔	↑
Time Anomalies		
Unusual time of day	✖	↓
Fast user movement	✔	↑
Simultaneous access from two distant locations	✔	↑

For our last example, suppose a hacker located in Asia has discovered the credentials to one of your corporate accounts. They use these to log in during your normal business hours, when one of your U.S. employees is also using these credentials. Adaptive Authentication would see this as:

The fact that the account is being accessed from two geographically remote locations would drive risk score higher. Even if the accounts were accessed a few hours apart, the fact that someone cannot fly from the U.S. to Asia in less than twelve hours would indicate abnormally fast user movement and also increase the risk score.

This elevated risk score would force both users—the employee and the hacker—to authenticate themselves via MFA. The hacker would fail to authenticate themselves and be blocked from accessing corporate apps.

Implementing Adaptive Authentication within Your Organization

Here are capabilities to look for when implementing Adaptive Authentication in your organization.

Built-in Analytics

It's critical to be able to view each login attempt, their risk scores, and the factors behind each score. OneLogin Adaptive Authentication provides exactly that, in realtime. For example, here is how the phishing attempt, scenario 1 above, would appear in to administrators using OneLogin:

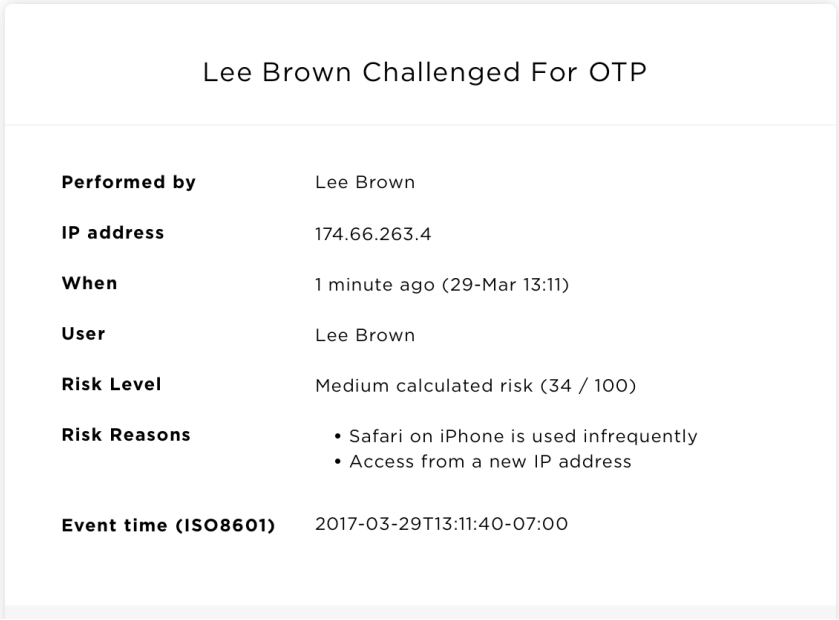

OneLogin Adaptive Integration integrates with your existing security infrastructure, including popular SIEM systems like Splunk, Sumo Logic, and Elastic; and popular MFA providers such as Duo, Google, and RSA.

SIEM Integration

Most security teams already have security analytics in place, whether from Splunk, Sumo Logic, Elastic, or other firms. OneLogin streams a range of authentication events to these systems. These events include login attempts and their risk scores, risk reasons, username, and IP address. Data is sent in JSON format, which can be consumed by a wide range of SIEM systems, and is sent in real time via streaming, as opposed to polling, for faster analysis of in-progress attacks. Once these events are in your SIEM system, they can be queried and analyzed like any other event.

Your MFA, made better

Some authentication products force you to use their MFA tool, forcing you into a closed ecosystem. Not OneLogin Adaptive Authentication.

In addition to working with our own OneLogin OTP, Adaptive Authentication also works with a range of MFA providers, including Duo, Google Authenticator, Symantec VIP Access, Yubico Yubikey, RSA SecurID, VASCO DIGIPASS and IDENTIKEY, Gemalto SafeNet Authentication Manager, and Swivel PINsafe. Authentications with high risk scores can be required to submit a second factor using any of these MFA products.

No user left behind

For users without smartphones, OneLogin Adaptive Authentication can send one-time passwords over SMS to provide an additional authentication factor. Security questions can also be used as an additional authentication factor. OneLogin comes with dozens of standard questions that are available in over 20 supported languages.

High Productivity MFA

Traditional multifactor authentication products have been a hassle to use: you have to unlock your phone, open up a mobile MFA, and push a “send number” button. Or worse, manually type in a number in a short amount of time. This friction has slowed the rollout of MFA that organizations desperately need to improve their security posture.

OneLogin Adaptive Authentication integrates seamlessly with OneLogin OTP, a mobile application that generates one-time passwords for use in multifactor authentication.

With OneLogin OTP, users see a push notification on their phone’s lock screen or Apple Watch, and can choose to Accept or Deny the login attempt. This provides a better user experience that increases the adoption of MFA to improve your organization’s security posture.

Multi Factor everything

Given the weakness of password-only authentication, it makes sense to apply MFA across all your IT assets, including SaaS applications, cloud infrastructure, VPN, and WiFi. OneLogin Adaptive Authentication secures SaaS apps such as Office 365, G Suite, and Salesforce.com; cloud infrastructure such as Amazon Web Services (AWS), Microsoft Azure, and Google Cloud; VPN from Cisco, F5, and Pulse Secure; and WiFi from Cisco Meraki and Aruba.

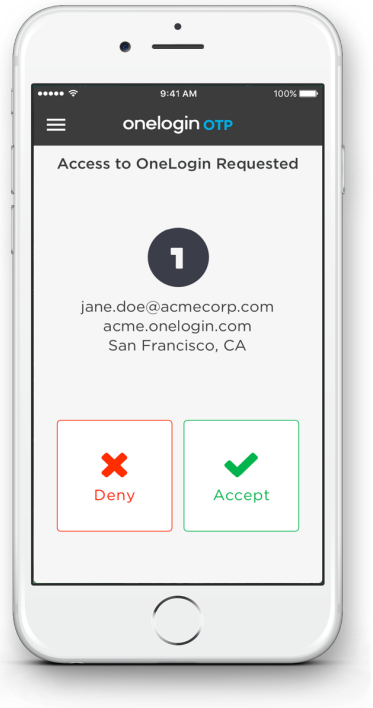

OneLogin OTP is a mobile application that generates one-time passwords for use in multifactor authentication.

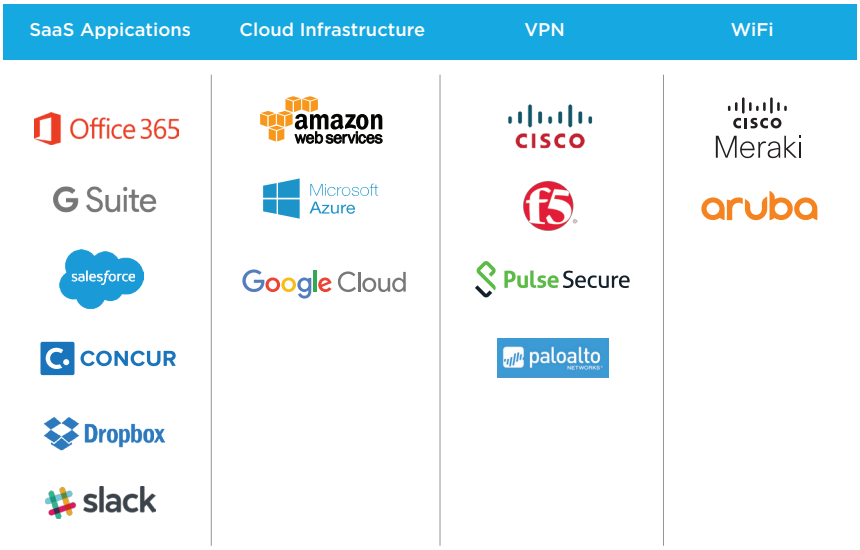

Get started today

OneLogin Adaptive Authentication is available today, and since it doesn't require complex static rules to be defined, it can be setup in just one minute. To start using machine learning to protect your organization from a broader range of security threats, contact us at onelogin.com/contact.

About OneLogin, Inc.

OneLogin brings speed and integrity to the modern enterprise with an award-winning SSO and identity management platform. Our portfolio of solutions secure connections across all users, all devices, and every application, helping enterprises drive new levels of business integrity and operational velocity across their entire app portfolios. The choice for innovators of all sizes such as Condé Nast, Pinterest and Steelcase, OneLogin manages and secures millions of identities around the globe.