opentext™ | Cybersecurity

Position Paper

The Next Generation of Access Control

As IT and security teams continue to struggle with the competing priorities of digital user engagement and security, a new, more powerful approach to identity and access management is needed. This paper outlines a complete approach that not only identifies the risk attributes of user/API access requests, but also measures the actual risk. We also review the fundamentals of risk assessment that are needed for automation to ensure that the right people have convenient access to sensitive resources.

Despite the fact that machine learning technologies have recently been introduced into adaptive access management solutions, they fall short in terms of measuring the risk of an access request. This paper describes the capabilities needed to automate access control.

Introduction

Even as the vast majority of business environments still rely on static entitlements to manage access to their resources, we know that the most secure environments have automated that process through access governance. The focus of this paper is the next step: executing those permissions in the form of *adaptive access control*, which is enforced at the point of and during the actual access request of protected resources, as well as follow-on requests. For more information about automating your access governance with am OpenText™ solution, please visit our identity and access management page at **www.microfocus.com/en-us/cyberres/identity-access-management**.

Assessing the Risk Level of the Identity Requesting Access

Now that remote access is no longer limited to road warriors and specialized offsite professions, there has never been a more important time for organizations to accurately assess the risk levels associated with access requests and identity claims. The recent pandemic has pushed teleworkers into the mainstream, vastly increasing the complexity of securing the resources that they consume.

Other trends that are increasing outsider threats include:

- Hybrid environments continuing their transition to the cloud.
- Organizations expanding their use of automation to share and consume sensitive data.
- Native mobile apps and the expanded use of automation to share or consume sensitive data.

And although identities tend to be static, their location, the device used, and other contextual information have the potential to change at any point of time. While this has always been true with workers who are frequent travelers, the massive expansion of teleworkers has increased the scale and complexity of this problem. The criteria most used to build a contextual profile includes elements such as recognized IP addresses, GPS location, browser cookies, and device IDs. Risk-based authentication integration kits such as the one offered in NetIQ Risk Service by OpenText™ allows for more specialized heuristics to be added as well.

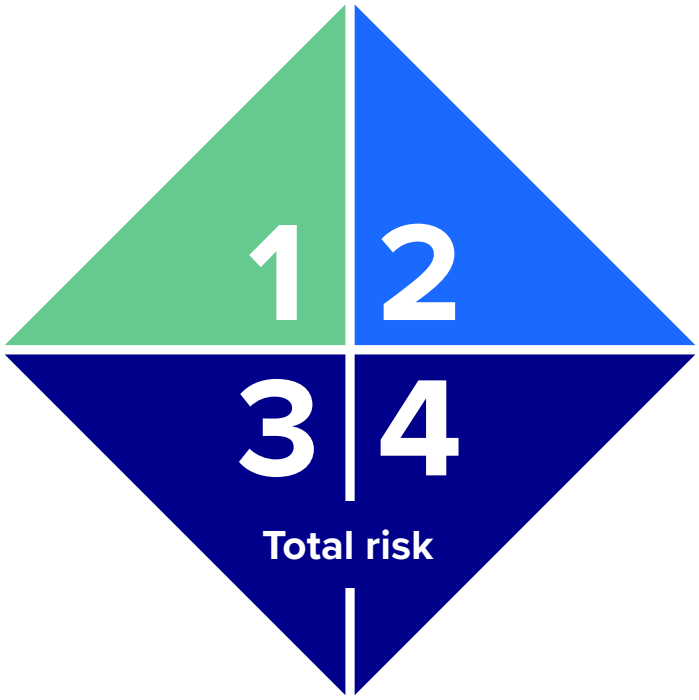

1

Richness of authentication types needed to verify identities across different situations

- Multi-factor authentication
- Strong authentication
- Rich set of methods
- Broad application and resource coverage
- Continuous authentication

2

Ability to respond to risk (control)

- Continuous authentication
- Session Management
  - Continuous authorization
  - End session

3

Measure risk that breached data imposes onto an organization (security-at-rest)

Potential metrics measured by an organization:

- Consulting and legal services
- Restitution
- Regulatory fines
- Recovery expenses

4

Ability to assess the risk level associated with an identity claim (security-in-motion)

- Device: known/unknown
- Location
- Behavior

Figure 1. While most risk engines can measure context and behavior, OpenText adds another dimension incorporating the risk of the resource itself.

Beyond defining risk-based policies, another type of technology gaining traction among IT and security groups is the application of machine learning technology to behavioral analytics. This type of analysis includes tracking behavior for the authenticated users' resource requests, such as:

- The types of resources consumed compared to the past.
- The speed and volume in which they are consumed.
- Regardless of defined rules, using past time/location/device patterns to categorize context as expected or unusual.

The reality is that regardless of assigned entitlements, over time, users show a pattern of actual usage that is a natural part of getting their job done. And as these captured patterns become more established and exceptions become more pronounced, they are a stronger indication of whether the identity is a false claim. It's also important to note that not all machine learning technologies are the same; they have varied capacity to identify characteristics that are meaningful to calculating risk accurately.

In general, it's helpful to break behavioral analytics into two categories:

- Real-time heuristics capable of measuring resource access behavior during each active session, in order to detect first-time cyber penetration.
- Comparing past behavior to recent session behavior, with a focus on performing a more in-depth analysis as a post-session exercise.

Given that most breaches of notable size take days, weeks, or even months for an outsider to set up before accessing large repositories of data, there is tangible value to using both methods to calculate behavioral analytics. Used together, they do more than enable IT teams simplify their access control rules: they identify vulnerability blind spots and become an essential component of a zero trust environment enforced at the application layer.

Measuring the Inherent Risk of Protected Data and Services

While user context and behavior are essential metrics for calculating risk at a point in time, one component that is often neglected is the risk inherent with the service or data itself. This means that, while security and IT teams have spent most of their time focusing on the context of the users to derive their risk scores, little if any effort has been spent defining the risk posed to the business by the crippling of a service or the breaching of a certain type of information.

Potential Impact of a Disabled Service

Although organizations don't typically include the impact of a disabled infrastructure or service to their business in the RBA risk scores, they usually have those calculations from the availability studies they perform as part of their infrastructure investment exercise. CIOs often ask themselves whether they can run their business depending on four 9's reliability, or do they really need to budget the extra money on five 9's? They make this decision based on analyzing the interdependencies and ramifications across their entire business for each service piece.

Inherent Cost of a Breach

Aside from the infrastructure impact studies, there are separate breach costs incurred that too often organizations don't bother calculating. While disabled costs are derived by calculating the price of an idle or partially idled business, the cost of a breach is far more diverse. Ponemon Institute offers the most complete breakdown and analysis of breach costs. Every business should conduct a review of their exposure using Ponomon's outline. Here are the categories of expenses that Ponemon Institute measured in their 2020 worldwide breach study*:

- **Detection and escalation**—Activities that enable a company to reasonably detect the breach (on average, this makes up 29% of total cost):
  - Forensic and investigative activities
  - Assessment and audit services
  - Crisis management
  - Communications to executives and boards

* **www.ponemon.org**

- **Notification**—Activities that enable the company to notify data subjects, data protection regulators, and other third parties (on average, this makes up 6% of total cost):
  - Emails, letters, outbound calls, or general notice to data subjects
  - Determination of regulatory requirements
  - Communication with regulators
  - Engagement of outside experts
- **Ex-post response**—Activities to help victims of a breach communicate with the company and redress activities to victims and regulators (on average, this makes up 26% of total cost):
  - Help desk and inbound communications
  - Credit monitoring and identity protection services
  - Issuing new accounts or credit cards
  - Legal expenditures
  - Product discounts
  - Regulatory fines
- **Lost business**—Activities that attempt to minimize the loss of customers, business disruption, and revenue losses (on average, this makes up 39% of total cost and lost business continues to be the largest single contributing factor):
  - Business disruption and revenue losses from system downtime
  - Cost of lost customers and acquiring new customers
  - Reputation losses and diminished good will

Mapping Cost to Actionable Risk Scores

One of the most effective ways to bring together the complete picture when calculating risk at a point in time is to take a pedantic approach to access governance. Access review and recertification campaigns are an essential part of an overall identity governance program. They should take into account not only the diverse personal persona, but also the incurred risk of the resources that they have been granted access to (per the analysis guidelines listed above). If both types of criteria are included, the resulting governance score can be leveraged as part of the total risk calculation.

Responding to Risk with the Right Access Control

Risk-based authentication (RBA) has been used for years to elevate authentication strength, typically invoking a second-factor authentication under pre-defined conditions, with user context being the most common approach. But the reality is that the risk level can change during a session:

- As discussed earlier, not all resources pose the same level of risk to the organization. This is true when a user accesses a higher risk resource during a session.
- If a session is hijacked, the risk level goes up immediately.

Both of these scenarios demonstrate the need to continue to assess the risk of each session until it is completed and is able to respond when it elevates to an unacceptable level. Session variables that need continual assessment are:

- Context of a user or programmatic request—assessing context at the beginning of a session protects against hacked credentials, continual assess protects against hijacked sessions.
- The user accesses a new resource that incurs a higher level of risk that may push the entire risk score beyond the threshold.

Continuous Authentication

Although security teams typically concentrate on beefing up their identify verification processes at the point when the user requests access to their resources, a more effective perspective is a layered approach—meaning that authentication isn't a one and done event, but rather a tool that is used as needed to manage risk by gaining greater confidence that the requestor's claim is indeed accurate. Consider the following scenarios.

LOWER RISK

Examples of low-risk access requests are general organizational information free from regulated or personal information, intellectual property, or financial data. Whether it's a username and password or a social credential that can be easily produced, the focus of these access use cases should be simplicity and convenience. Continuous authentication doesn't offer value in these situations.

HIGHER RISK

Now that virtually everyone and everything is digitally connected, the current standard for protecting sensitive and regulated data requires a two-factor authentication or some other type of strong authentication. Typically, authentication is invoked when access to digital resources is being requested. There are two basic approaches that IT teams are taking.

- The most common approach to verifying an identity claim is to require a password at the time of request (user or programmatic), followed by a request for a second factor of authentication or even another instance of the same factor. This approach is quite effective in protecting against credentials that have been compromised through either phishing or a hacked repository. Government policies protecting regulated data require two-factor authentication.
- Increasingly, organizations are turning to passwordless authentication as the primary type for protecting their resources. Not only does this reduce the number of passwords that an individual needs to remember, but it's usually more convenient for the user while still being highly resistant to phishing. Fingerprint, facial, and FIDO devices are currently the most common passwordless technologies, but they also have their limitations. Cost will continue to be the biggest barrier to adoption, but it's important to note that as long as security teams provide a password-based backup to the passwordless method, it will continue to be vulnerable to phishing attacks.

RAISED RISK

Despite increasingly stringent security regulations, persistent breach rates illustrate why a new approach to access control is needed. The key attribute of continuous authentication is that identity verification is not a one-time event. As such, there are two distinct advantages:

- Context and behavior can continue to be monitored for signs of impersonation throughout the session.
- In response to a rise in the risk score (context, behavioral indicators, or access of higher risk information), the identity can be re-verified through another authentication. Through both continued monitoring of the session and the ability to re-verify the requester's identity, you can invoke an identity verification request not only at the beginning of a session, but also throughout when the user/API submits additional access requests to the same or a different resource.

Continuous Authorization

While continuous *authentication* is the ability to verify an identity throughout the entire session, continuous *authorization* is the ability to control access based on the current assessed risk. Situations where authorization might need to be adjusted during a session include:

- A context is triggered that is outside corporate policy.
- Behavioral analytics spikes the risk score.
- The consumer fails an authentication invoked during the session.

Organizations will likely define other continuous authorization scenarios based on specific needs. And since prescribing too many rules can quickly become overwhelming, the best strategy is to keep them simple and few in number. This approach will also help protect against human error. One important rule is the ability to not only deny initial access in high-risk situations, but to also terminate a session if a behavioral or contextual risk-related threshold is met.

Why You Need a Rich Set of Authentication Methods

The power of continuous authentication to invoke a method at any point during a session introduces new requirements into the user experience. Rather than being a one-time event potentially for the entire day, typically through single sign-on (SSO), continuous authentication is the antithesis of leveraging a single method to access everything. The key question becomes, how do organizations make such a dramatic shift to security without making the environment unusable? While here are several approaches, there are likely others that might provide a better fit.

Passive

Examples of passive authentication types are facial recognition, typing speed, or mouse movement. Mobile gestures such as force of touch and other types of gesture movement are also used as indicators.

Multiple Passive Options
In a paradigm of employing multiple authentications throughout a single session across a range of situations, the more options an organization is able to deploy, the more effective they will be in inconveniencing their users. This is best done through an authentication framework that is robust, supports open standards, and aggressively adopts new methods as they are available.

Low Friction
In many situations, a low-friction option works well for users. For example, employees with specialized responsibilities are fine with being prompted to touch a fingerprint reader or having to carry their smartcard, FIDO, or Bluetooth-enabled device with them—all of which are passwordless. In these situations, passwordless is key because the authentication isn't something that has to be remembered or typed in. A simple touch, or less, has the user on their way.

Authentication Disruption
Based on the sensitivity of the resource being accessed and measuring what's at stake for the organization and assessing the risk, there are situations where a strong authentication request is warranted in order to enforce the appropriate level of security. It might be a last-ditch option to allow the user to continue a session before terminating it. However, organizations that make disruptive authentication requests the norm as part of their zero trust or continuous authentication implementation will see lower productivity and engagement and likely user revolt.

Shopping for an Authentication Framework That Can Deliver
Retaining and maintaining control of sessions where valuable information is being accessed presents a substantially more rigorous set of requirements. In fact, the most fundamental requirement of the framework is that it's capable enough for the organization to consolidate their authentication silos onto it:

- **High performance**—Both the authentication request and the response to it (access) need to be fast, especially in consumer scenarios.
- **Morph to the shape of the organization**—Some are centralized, while others are highly distributed. Ether way, or somewhere in between, the framework needs to be scalable and responsive.
- **Standards-based**—Easier integrations and no vendor lock-in.
- **Broad application and platform support**—Beyond just web, iOS, Android, Windows, OS X, and Linux.
- **Diverse set of methods**—Smartphone, Windows hello, Voice OTP, facial, SAML, PKI PKCS-11 RFID, PIN code, geo-fencing, Windows Hello for Business, Microsoft Live, hard tokens, fingerprint, OAuth2, PKI PKCS-7, NFC, Bluetooth, Google Auth, SMS OTP, TouchID, Windows OTP Tool, Mac OSX OTP Tool, challenge response.

The more complete your authentication framework, the deeper and more securely you'll be able to engage with your consumers and partners. It provides more opportunity to verify someone's identity before restricting or terminating their access. It's an essential component for any organization trying to achieve a zero trust environment.

Summary

There are three areas where organizations can achieve deeper, more effective interaction with services and digital users:

- While risk-based authentication has become commonplace, incorporating varied levels of risk inherent with different information types as part of that calculation is rare. Measuring the inherent risk of each type of sensitive or regulated data enables a far more accurate risk score.
- Continual access management of a session (user or programmatic) is the foundation of the next generation of identity and access management.
  - Rather than simply evaluating the initial risk of an access request, the continual monitoring of a session for changes in the measure criteria (context, behavior, risk of information types) offers a significantly more thorough approach to driving protective measures.
  - Continual authentication is a more exhaustive approach to identity verification. It's a potential reaction when a risk score peaks beyond a threshold during a session.
  - Continual authorization is the ability to restrict access to a narrower set of information or to break the session altogether when the risk of an access request within a session reaches unacceptable levels.
- The larger an organization's library of authentication types, the wider variety of situations they are able to accommodate—especially if that library includes a variety of passive authentication types that don't interrupt the user. The larger the number of lower-friction authentication methods available for use in an organization, the better able they are to provide a faster and simpler user experience.

The NetIQ portfolio by OpenText offers the most complete identity and access management platform and is tailored to the most complex environments.

Learn more at
www.microfocus.com/en-us/cyberres/identity-access-management

www.youtube.com/NetIQUnplugged

About NetIQ by OpenText

OpenText has completed the purchase of Micro Focus, including CyberRes. Our combined expertise expands our security offerings to help customers protect sensitive information by automating privilege and access control to ensure appropriate access to applications, data, and resources. NetIQ Identity and Access Management is part of OpenText Cybersecurity, which provides comprehensive security solutions for companies and partners of all sizes.

Connect with Us

www.opentext.com

opentext™ | Cybersecurity

OpenText Cybersecurity provides comprehensive security solutions for companies and partners of all sizes. From prevention, detection and response to recovery, investigation and compliance, our unified end-to-end platform helps customers build cyber resilience via a holistic security portfolio. Powered by actionable insights from our real-time and contextual threat intelligence, OpenText Cybersecurity customers benefit from high efficacy products, a compliant experience and simplified security to help manage business risk.