1KOSMOS

The Promise of Biometric Authentication Versus the Threat of Deepfakes

Modernizing Identity Access Management to improve both user convenience and security

Whitepaper

Executive Summary

In recent years, the rise of deepfake technology has introduced a new and concerning threat to online security. Deepfake videos, images, and voices, powered by sophisticated artificial intelligence algorithms, allow cybercriminals to create convincingly realistic impersonations of individuals, including their facial expressions and voices.

Bad enough that these videos can fool humans, but worse, as organizations turn to biometric authentication to replace passwords, deepfakes expose vulnerabilities that many first-generation biometric authentication systems had not anticipated and cannot address.

This whitepaper delves into the intricacies of deepfake attacks, exploring their types, tactics used for identity impersonation, and the role of biometric authentication with liveness detection and injection attack detection in mitigating these threats.

By understanding the threats and risks presented by deepfake attacks and implementing appropriate security measures, Information Technology and Security executives can modernize identity and access management to improve the user experience and better safeguard their organization while protecting their workers' and customers' online accounts from compromise.

Introduction

In 1997 when IBM's Deep Blue beat grandmaster Garry Kasparov in a game of chess, Machine Learning entered the modern lexicon. A decade later and enabled by progressively more powerful chip sets, machine learning evolved into deep learning, including layered, hierarchical logic that equipped computers with the ability to recognize complex patterns buried in large datasets, draw inferences, and self-adjust without manual intervention. Deep learning then found its way into the hearts and minds of consumers embedded in Siri, Alexa, and so many other virtual assistants as they listened and responded dutifully to voice commands.

Fast forward to 2014 and learning algorithms move beyond pattern recognition and predictions. Ian Goodfellow has invented the Generative Adversarial Network (GAN) pitting one deep learning neural network against another in a challenge to produce a fake image from statistical analysis of an original. The two networks work as adversaries in an iterative cycle to produce, grade and refine the image for accuracy. Iterations follow at the speed of light. Generative AI is born. Computers now create digital media without following explicit instructions.

Today, generative AI is in the movies and television commercials we watch and in the computer games we play. It is driving cars on city streets, predicting protein structures and folding (e.g., AlphaFold), accelerating materials design (e.g., decomposable plastics) and helping clinicians analyze CAT scans to name just a few commercial applications. It is no wonder criminal elements are using it to attack the places where we work, family and friends we hold dear, and ultimately our identity, assets, and livelihoods.

But unlike 2014, today you do not need to be one of the brightest minds in data science to create a deepfake. Less than a decade later, the ability to create deepfake content has been packaged for nontechnical users. Inspired to advance learning, creative expression, and entertainment, but repurposed for far less noble intentions by political extremists, social media trolls, cybercriminals, and others, this new generation of synthetic media blurs the lines between fantasy and reality. It is a fool's game to ignore the threats and the risks.

Anything that can be digitally sampled can be deepfaked … an image, video, audio, even texting to mimic the style and syntax of the sender. Equipped with any one of a half dozen or so widely available tools and a training dataset (e.g., YouTube videos), even an amateur can produce deepfakes of sufficiently quality to coax **one finance worker out of $25m** or **swindle a few thousand bucks** from an unsuspecting grandparent.

At the crux of it, what makes deepfakes so effective … a few million years or so of human evolution that has programmed each of us to trust and believe what we see and hear. More so than any cyber threat before, deepfakes attack human sensibilities, inspiring fear, affecting our decision making and causing us to act in ways that seem valid, but are inspired by a ruse.

For human decision-making, a measure of awareness, education and street smarts combined with governance and oversight in the corporate domain can help mitigate the threats. But, what about deepfakes turned against biometric authentication?

It is no coincidence that deepfakes are surging at a time when organizations are moving in mass to passwordless authentication and away from antiquated, message-based multi factor authentication (MFA). The "who you are" or inherence factor that biometrics represent is replacing the "what you know" in the form of passwords, one-time codes, or personal factoids, such as the make and model of your first car.

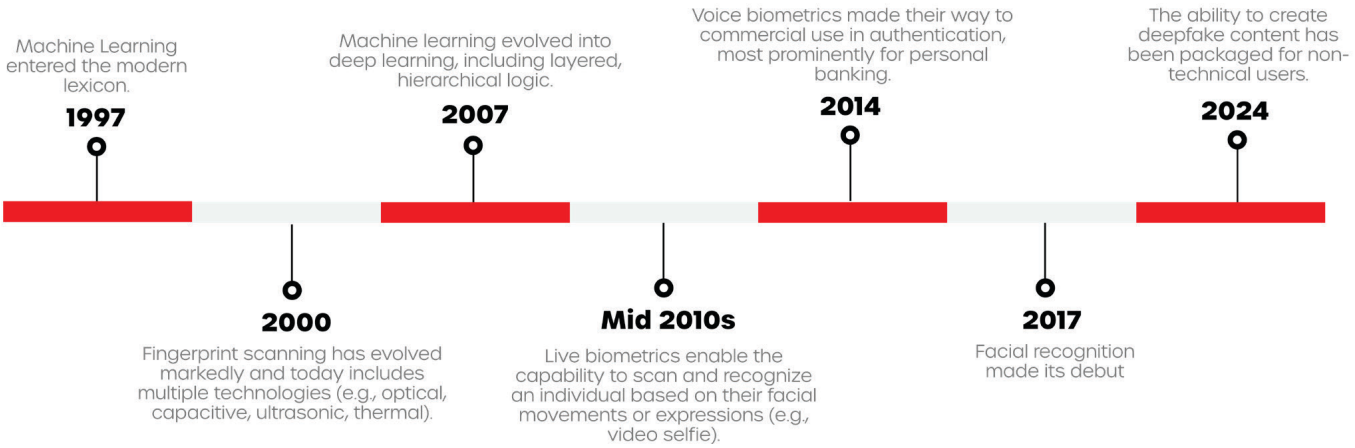

Biometric authentication fundamentally alters the traditional tradeoff between security and convenience by improving both, but they usher into the forefront many inescapable questions about efficacy and security. Are they safer than passwords? What new risks do they introduce? What about privacy? And, which methods are more trustworthy than others?

Let's look at the various types of biometrics, their benefits, limitations, and the attack vectors exploited by deepfakes. But first, just a bit of context on the types of deepfake attacks.

Types of Attacks

Deepfake attacks come in two varieties:

One.

Presentation Attacks

Two.

Injection Attacks

Presentation Attacks

Presentation attacks involve presenting a fake image, rendering or video to a camera or sensor for authentication. Examples include:

Print attacks:

- 2D image
- 2D paper mask with eyes cut out
- Photo displayed on smart phone
- 3D layered mask
- Replay attack of a captured video of the legitimate user

Deepfake attacks:

- Face swapping
- Lip syncing
- Voice cloning
- Gesture / expression transfer
- Text-to-speech

Injection Attacks

Injection attacks involve manipulating the data stream or communication channel between the camera or scanner and the authentication system. Such is common with a virtual device or man-in-the-middle (MITM) attack. Using automated software intended for application testing, a cybercriminal with access to an open device can inject a passing fingerprint or face ID into the authentication process, bypassing security measures and gaining unauthorized access to online services.

Examples Include:

- Uploading synthetic media
- Streaming media through a virtual device (e.g., cameras)
- Manipulating data between a web browser and server (i.e., man in the middle)

Types of Biometric Authentication

Biometric markers for use in authentication comprise a virtually endless list of possibilities including retina, iris, hand/palm, ear, vein, heartbeat, signature, gate, DNA and even odor. For our purposes, we will focus on the few that comprise the broadest use and widest commercial adoption:

One.	Two.	Three.
Device Biometrics	Voice Biometrics	Live Biometrics

Device Biometrics (e.g., Touch ID, Face ID)

Fingerprint scanning has evolved markedly from its first consumer introduction in 2000 and today includes multiple technologies (e.g., optical, capacitive, ultrasonic, thermal). As an example, Apple's "Touch ID" uses a combination of capacitive and ultrasonic imagery to provide a match probability of 1 in 50,000.

Facial recognition made its debut about a decade later and in 2017 increased match probability to 1 in 1,000,000 by combining machine learning with depth sensing cameras that can capture multiple 3D infrared images from various angles to model the face.

Most leading device-level systems do not store biometric images but rely on mathematical representations. The data supporting those models is encrypted in communication with the OS and while in storage in the Secure Enclave of the device. This makes compromising the device biometric for presentation attacks difficult at best.

Fingerprints do not change with age and offer the convenience of logging in with a simple touch. Faces do change over time, but machine learning can discern typical changes making an even more convenient user experience.

While injection vulnerabilities on open devices have been reported in various Android devices, injection attacks targeting device biometrics are rare and require physical access to a device.

Device biometrics address user convenience, but they fall short of high identity assurance security requirements, such as accessing a corporate VPN, privileged systems, or high value consumer purchases. This is because with device-level biometrics enrollment is informal.

Specifically, anybody who gains administrative access to a device can register their biometric. This could include family, friends, pranksters as well as those with criminal intent. So, while a biometric scan can match a biometric previously registered on the device, the identity associated with that biometric is not established with high assurance. Whether device-level biometrics offer sufficient security is conditioned on the risk associated with granting systems access.

Voice Biometrics

After slow baking over a 40-year timespan, voice recognition began widespread adoption in the early 2000s following advancements in machine learning. A little over 15 years later, voice biometrics made their way to commercial use in authentication, most prominently for personal banking.

Like device-level biometrics, authentication via voice biometrics requires a real-time match against a reference sample captured during registration. Similarly, the typical voice biometrics system does not capture a voice recording, but rather a "voiceprint" comprised of a mathematical model, encrypted, and securely stored. The modeling is sufficiently sophisticated to match the way an individual pronounces words, and includes various other unique factors including tonality, intonations, cadence, and pitch.

To the human ear, deepfake audio can be exceedingly difficult to detect.

Match probability for voice depends on several factors including voice quality, but it is highly accurate though slightly less user friendly than device biometrics because the user may have to repeat a specific phrase captured at registration.

Registration of voice biometrics shares some of the same shortcomings as device biometrics. Any person able to recite accurate personal details can register and then use voice biometrics. This again addresses user convenience but offers little security beyond knowledge-based factors.

For presentation attacks, voice conversion and text-to-speech both use generative AI trained on a voice sample, for example, from an online video or voice recording. The synthetic media can then be presented to the audio mic in the device. Technically, injection attacks are also possible, but it is relatively easy to detect a virtual device.

To the human ear, deepfake audio can be exceedingly difficult to detect. Sophisticated defenses using generative AI can detect slight variations in pitch, speech patterns, and choice of words. Statistical analysis of audio spectral features and even computer vision modeling of vocal pitch, timbre, the dynamic range of voice wavelengths and other inconsistencies have also been refined. And, in an approach reminiscent of antivirus solutions, AI defenses improve when fed modelling data from new fakes as they appear, which makes continuous monitoring and sampling essential for AI defenses to resemble actual learning.

Extending Passwordless to Restricted Environments

The 1Kosmos platform provides passwordless deployments in diverse environments like call centers, manufacturing floors, kiosks, clean rooms, SCIFs (Sensitive Compartmented Information Facilities), kiosks, shared workstations, and healthcare facilities where multiple users may login to the same device.

These use cases do not allow mobile devices, preventing organizations from deploying Microsoft Authenticator, and due to the limitation of the number of enrolled users per device and lack of support for desk "hoteling," Windows Hello for Business can be rendered unsuitable.

The 1Kosmos 1Key has a fingerprint reader and supports one-to-one, one-tomany, and many-to-one use cases. For example, one security key can support multiple users logging into a device with either separate or shared accounts and still retain an immutable record identifying a specific individual.

The approach is simple – install the 1Kosmos 1Key on the protected device, providing a biometric passwordless login for any user onboarded within the control plane. Users register once and can then authenticate on any device.

IT and security teams can continue to leverage conditional access policies in Entra ID while leaving the authentication to 1Kosmos. This approach also contains costs because users do not need their own keys, so fewer keys and replacements are needed. It also prevents security vulnerability from unauthorized key sharing

By integrating the 1Kosmos 1Key Biometric Security Key into Entra ID, organizations can:

- Deliver passwordless authentication to restricted areas
- Reduce the risk of password-related vulnerabilities
- Utilize biometric authentication with high identity assurance
- Support MFA login by default with a FIDO-compliant authentication method
- Reduce cost with a register-once-use-anywhere, workstationindependent strategy
- Simplify management as users will no longer be required to carry an assigned authenticator

1Kosmos as an Entra ID External Authentication Method (EAM)

Microsoft's External Authentication Method (EAM) is a new feature in Microsoft Entra ID that allows organizations to satisfy multifactor authentication (MFA) requirements using external providers, and not rely on native Microsoft authenticators. Organizations can deploy 1Kosmos and, therefore, utilize verified identity and passwordless multi-factor authentication (MFA) for every login.

The 1Kosmos platform offers flexible authentication options with 11 different authentication methods through which organizations can match conditional access policies and tailor to risk levels.

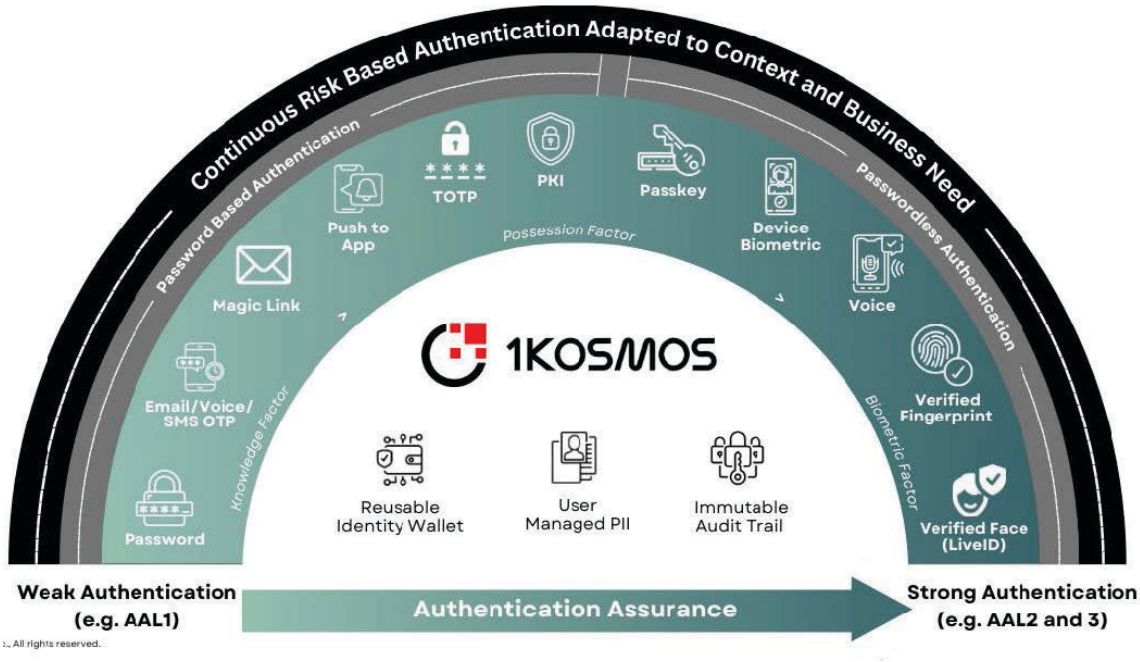

These same methods can be deployed to environments not natively supported by Entra ID, such as legacy on-premises technologies and various operating systems, ensuring a consistent login experience no matter the platform, application, or service. This will also benefit security teams as this integration will significantly reduce the number of authenticators required for day-to-day business.

By integrating the 1Kosmos as an EAM to Entra ID, organizations can:

- Improve user experience across systems
- Continue to leverage conditional access policies
- Match authenticator to activity risk
- Reduce management cost by eliminating additional authenticators

Interoperability

Interoperability and extensibility are key features of 1Kosmos, which gives organizations using Entra ID a safe and secure way to get new users and new organizations rapid access to digital services. Out of the box, 1Kosmos comes with over 50 connectors, an open API framework, and a flexible SDK.

This allows for rapid integration with Entra ID and all other technologies and legacy systems that fall out of scope for Entra ID passwordless, plus easy embedding of 1Kosmos into mobile apps.

Our solution also provides authentication resiliency and interoperability, where organizations can deploy 1Kosmos as the primary authenticator and leverage Microsoft Authenticator as a backup.

By leveraging the 1Kosmos interoperability capabilities in conjunction with Entra ID, organizations can:

- Improve coverage across a wide variety of applications and services
- Swap underlying technologies without impacting the authentication user experience
- Improve user experience with a consistent passwordless experience
- Provide a consistent, passwordless authentication experience across systems
- Continue to leverage conditional access policies

Conclusion

The integration of 1Kosmos with Microsoft Entra ID offers organizations a solution to prevent identity-based attacks while facilitating convenient passwordless authentication across disparate IT systems. As discussed, the combination of 1Kosmos and Entra ID addresses critical security gaps, preserving the Entra ID investment while improving user experience across diverse environments.

1Kosmos + Microsoft Entra ID Highlights:

- Streamlined self-service user onboarding and identity verification: reducing administrative overhead, so security teams can leverage verified identities for authentication, password resets, and account recovery.
- Extended passwordless authentication beyond Microsoft Entra ID: covering legacy systems, various operating systems, and critical infrastructure – VPN, PAM, virtual machines, and more.
- Secure password reset and account recovery options using live biometrics: minimizing IT support costs while providing user convenience and maintaining high identity assurance.
- Provide passwordless access for edge use cases: extending passwordless to eliminate identity-based attacks for one-to-many and many-to-one use cases.

- Flexible authentication methods: tailored to risk levels, ensuring consistent login experiences across all platforms and applications with minimal transaction friction.

- Improved interoperability: with systems through numerous out-of-the-box connectors and industry-standard protocols including 1Kosmos as an EAM via WS Fed and WS-Trust.

This integration delivers high identity assurance, self-service identity verification, and a passwordless approach to cover all use cases. The result is a safe and efficient way to onboard, authenticate, and verify new and existing users for access to digital services, password resets, and account recovery requests.

As cyber threats continue to evolve, this combined approach offers a flexible forward-thinking solution that improves the balance between security requirements and user convenience, making the combination an invaluable asset for security teams looking to eliminate identity-based attacks and minimize non-value-added administrative overhead related to user onboarding and account / password recovery.

About 1Kosmos
1Kosmos enables remote identity verification and passwordless multi-factor authentication for workers, customers and residents to securely transact with digital services. By unifying identity proofing, credential verification and strong authentication, the 1Kosmos platform prevents identity impersonation, account takeover and fraud while delivering frictionless user experiences and preserving the privacy of users' personal information. 1Kosmos performs millions of authentications daily for government agencies and some of the largest banks, telecommunications, higher education, and healthcare organizations in the world.

1KOSMOS