JUNE 2023

UNLOCKING THE FUTURE:

Digital Identities in an AI World

Authored By:
Jim Routh, ICIT Fellow

ICIT Institute for Critical Infrastructure Technology
The Cybersecurity Think Tank

Unlocking the Future: Digital Identities in an AI World

June 2023

Authored by Jim Routh, ICIT Fellow

Contents

Introduction

Identity & Access Management (IAM) is often called the "plumbing" of every enterprise. However, it has recently evolved into a "bastion of technical innovation" built on a data science foundation, dramatically changing how cyber security controls can be applied within an enterprise. To describe how different this is, I reflect on my experience as a cybersecurity leader in seven industry-leading enterprises in the past two decades. I was told that IAM was the "dumping ground" of mediocre IT talent, a place for those who struggle with the constant change and technical challenges of enterprise IT management. More often than not, IAM was part of the cybersecurity function/organization. It represented repetitive administrative work that was labor-intensive and somewhat cumbersome due to outdated legacy systems.

This has changed today; technically oriented professionals consider working in IAM cool. This is due to the maturity of AI/ML models available in software products and within reach for those organizations that invest in advanced data analytics. Using AI for IAM controls is very cool for technologists (as most things in AI are). IAM controls designed with AI enable authorities to work in real time without human involvement. This improves the user experience, lowers operating costs, and manages risk more effectively. This is not your grandfather's Oldsmobile!

When an employee joins an enterprise, the leader (or HR) requests network access that follows an established workflow where someone in the IAM function grants the approval and provides the employee with a user ID and a password. Depending on the role, the employee may be issued certain entitlements because of the job family. These entitlements are also approved by someone in IAM in a specific workflow in addition to additional approval from the data owner.

Access approvals, provisioning, annual recertifications, and de-provisioning require an administrative workflow designed by IAM professionals for enterprise-wide applicability. Many of these workflows originated with legacy systems decades ago. At one of the enterprises where I worked, we counted all access recertification requests along with the wait time embedded in the standard workflow for approvals. We discovered that the wait time for leaders represented over 300 years of total time over one year! We also calculated that more than 15% of a manager's time annually was applied to access recertifications driven by regulatory compliance needs. I learned that the perception that weaker technical talent was attracted to the repetitive administrative work within the IAM function was not far off the mark.

In today's enterprise, it's possible that the systems provisioned for many end users are cloud-based (SaaS) and no longer exclusively legacy systems accessing data stored in proprietary data centers. Thanks to data science advances, capturing online behaviors across attributes for specific end users and establishing a behavioral pattern for each user is relatively straightforward. The pattern and corresponding risk score based on the pattern can automatically be used to provision real-time entitlements without human approval in a workflow. This eliminates the need for all provision requests to be approved by multiple leaders in a workflow. In most cases, 80% of the approvals for provisioning entitlements can be eliminated. This represents a significant improvement in fulfillment time, improving productivity for employees.

The Role of Artificial Intelligence and Machine Learning

Regulatory requirements often mandate entitlement recertification on an annual basis. Most recertifications required to meet regulatory requirements can be automated using AI/ML models. In this case, the user provisioned with a specific entitlement has a behavioral profile based on their normal online activity. The recertification can be validated and approved without human review if the risk profile is within an acceptable deviation range. Using a behavioral profile, this approach eliminates the need for humans to research and approve recertification requests for most employees and represents a significant improvement in productivity for the enterprise. No more long wait times for recertification approvals. Regulators will recognize that this is more effective than relying on managers to approve recertification requests that they often don't take the time to understand. In addition, data analysis of specific entitlements provisioned and used enables an enterprise to eliminate entitlements not used to enforce the objective of least privilege. This results in lower operating costs, better user experience, higher productivity for leaders, and better enterprise risk management. This makes for a compelling business case to transform the IAM function and redesign processes using data science (AI/ML models) as a foundation.

User Behavioral Analytics (UEBA)

User behavioral analytics (UEBA) plays a significant role in the IAM function today and as we advance. This does not exclusively require a UEBA software product or a specific type of IAM platform. The reality today is that there are several ways to use behavioral analytics that may or may not include:

- An IAM platform
- A UEBA platform
- A passwordless authentication platform
- A data lake management platform
- A DLP solution
- A fraud operations platform
- A threat detection solution

The key for any enterprise is to acquire and develop IAM talent with data science knowledge and a commitment to understanding model-driven security control design. The good news is that many technology choices incorporate AI/ML functionality that can be deployed for IAM needs for both on-premise and cloud-based applications.

How does an enterprise start this transformative process, and how can an existing IAM function justify the business case for this level of change? Given the constraints of current market conditions, these are questions often posed by CISOs that must deliver favorable business results from any cybersecurity investments. The answer is that the investment required for an IAM transformation should be based on a business case of lower operating costs, higher productivity for managers, and significantly improved

service levels for employees. A byproduct of the IAM transformation is more effective cyber resilience for the enterprise.

Lower operating costs

Existing IAM processes are administrative and labor-intensive. Using data science skills to automate manual approvals reduces the headcount required in legacy processes designed for human administrative decisions. You can decide how best to eliminate roles while creating new roles based on different skills (data science and data analytics). The lower headcount results in lower operating costs on an annualized basis. A reasonable expectation is to reduce the headcount requirement by 25% due to the automation of key processes. There is a need to increase compensation levels for the data science skillset (data scientists have higher costs vs. IAM administrators), so the 25% annualized cost decrease will be offset by the increase in compensation levels for the new skill set. Therefore, the net impact may be closer to a 15% overall cost reduction vs. 25%, depending on the added talent acquisition costs. I observed a staff turnover in IAM of 60% over a two-year cycle in several enterprises; however, this included IAM staff movement to new roles based on their professional development desire and commitment.

Increased Productivity

As you upgrade workflows by applying advanced data analytics to reduce the requirement for humans to make decisions significantly, there is a clear reduction in time for decision-makers representing a productivity gain. Measuring the growth in business value applied to the increase in capacity for managers is difficult. Still, it should be relatively easy to determine the elimination of approvals and the average time it takes to review and approve a recertification approval transaction. Encourage your newly minted data science professional to help you complete the analysis. Productivity gains of 10-25% are normal when redesigning access recertifications, specifically applied to leaders.

Improved service levels

Using advanced data analytical models to drive transactions dramatically reduces the wait time for approvals for most transactions, resulting in substantial improvements in wait time. Most network access requests for new employees based on job roles can be processed in hours vs. days and then immediately revoked if the background check fails. Most privilege entitlement requests can be automatically approved (based on risk profiles without human review) while applying real-time monitoring of online behavior for the time the entitlement is used. All of these changes translate into much higher service level responses for enterprise users, which also improves productivity for the enterprise. Committing to a 50% increase in service levels on average for the redesigned processes is a reasonable commitment to make. The actual results will likely be higher.

IAM Belongs With CISOs

To accomplish the transformational effort, a CISO should own the IAM operations during the project timeline (18-24 months) to retool and redesign the process. It also enables the CISO to encourage a commitment to professional development for IAM staff. Afterward, moving IAM operations into the IT operations function is feasible if desired. Like any transformative project, developing a talent strategy to support the new IAM model is key. One successful approach is to divide the IAM team into Plan, Build,

and Run functional areas based on what they wish to learn (skills), allowing them to choose. The Plan team develops the blueprint and implementation plan identifying IAM platforms to retain, revise or de-commission over time. The design team creates the new workflows and focuses on the advanced analytics needed to use models and streaming data. Often this requires a data lake environment dedicated to the IAM team or cybersecurity. I have used the enterprise data lake for IAM and built a dedicated data management capability. The Plan team can evaluate and decide on third-party capabilities to support advanced data analytics. Ensure you recruit a few resources with data science experience to the plan and design teams. Think of data science talent supporting two dimensions of work effort:

1. Improving data quality for things like key performance indicators
2. Designing and implementing AI/ML models applied to streaming data

Both contributions from resources with data science expertise are essential for the transformation process. Don't worry if your data scientists have limited cyber experience; they will likely be hungry to learn domain expertise in cybersecurity over time. IAM professionals that understand the fundamentals of applying data science (AI/ML models) to IAM transactional data represent the skillset that you should emphasize as most useful to the IAM team.

Data Science Fundamentals

IAM professionals need to understand how to use data analytics to design and implement IAM controls that are not dependent on the consistent use of administrative labor. The fundamental concepts are as follows:

1. ***Behavior does not lie.*** The FBI learned this when interrogating potential criminals and reading micro expressions or body language when subjects answered questions. It applies in IAM to online or accesses behavioral patterns. Deviation of behavioral patterns provides an opportunity to initiate an automated workflow for risk management
2. ***Cluster analysis to determine user/customer online behavioral patterns*** and then identifying deviation from patterns has a lower degree of difficulty than discovering behavioral patterns of threat actors. Threat actors change behaviors to crack systems. End users typically follow patterns that are easily identified.
3. ***Many relatively benign online behavior attributes*** **can** be identified to establish a baseline pattern. Measuring the deviation of an established pattern is straightforward. For example, suppose a user typically logs into the network and uses a calendar management tool every business day at 8:35 AM. In that case, this attribute (time) represents a pattern. Deviation [using a specific measure] from this and other attributes can trigger an additive authentication or verification process without labor from the IAM team.
4. ***Patterns across multiple attributes can be used to establish a baseline risk score****.* End users that access the same information daily follow a set pattern and a risk score based on the pattern. The risk score can automatically provision additive entitlements and recertify the same without requiring IAM staff to provide approvals in the workflow.

5. ***Data analytics from transactional data sources is essential for increasing automation and decreasing labor time in IAM.*** Data attributes are readily available in transaction processing systems and specifically IAM platforms. These attributes can be applied effectively to increase automation through the use of models improving service levels at a lower operating cost

6. ***Measuring the deviation of the pattern requires using a threshold (specific number)*** to trigger an automated workflow action. For example, if variation across attributes is between 5.0 and higher, this could trigger a privilege suspension. If the score is 3.0-4.9, this could trigger a flag/alert for follow-up. No actions are required if the score is 2.9 or below since this is within the acceptable deviation range. Using specific numbers (threshold scores) gives the enterprise great precision in adjusting the deviation ranges.

7. ***Use data science staff to contribute to the evaluation of software platforms.*** More and more IAM platforms and tools are incorporating AI/ML into their platform's advanced analytics. Determining the level of AI/ML knowledge/skill in the vendor's platform is better suited to data science professionals, so invite them into the vendor evaluation process for any IAM capability.

8. ***Encourage IAM staff to invest their time in learning new skills.*** Cybercriminals change tactics consistently based on the success and failures they encounter. Therefore, cybersecurity professionals must embrace the need for continuous learning and adjusting controls based on the change in threat actor tactics. A good way to measure the commitment to making adjustments is to encourage all IAM staff to identify the skills they wish to invest their time in mastering as a part of a professional development plan. This can form the basis for a dialog with their respective leader and the CISO. Leaders can encourage IAM staff to invest in data science fundamentals applied to IAM control design to improve their marketability (professional choices).

These principles form a new foundation for IAM professionals and require a different level of technical and analytical skill. Many IAM staff that thrived in an environment of labor-intensive, repetitive administrative tasks and approvals may reject the need to retool analytic skills or be at a point where significant investment in learning new skills is not feasible. Your role as a leader is to serve employees and those unable to make a substantial investment in learning new skills and encourage them to settle for areas in the enterprise where they can contribute. That opportunity may have historically been in IAM, but that may not be true. This transition is often difficult for leaders and employees, given established norms in IAM over the years. However, new skills and advanced analytics are essential for IAM to transform. You will discover that a redesigned process with transactions using models to trigger automation requires far fewer resources performing administrative tasks which means fewer people and ultimately lower operating costs.

IAM Platform Requirements

Conventional IAM platforms are designed to increase transactional processing with improved record-keeping capability to manage access risk while meeting compliance requirements for on-premise or legacy systems. These requirements are representative of legacy IT infrastructure.

IAM platforms should be designed to support cloud-hosted applications and those still running enterprise data centers. The IAM platform should enable the goal of Zero-Trust and least privilege. More specifically, advanced data analytics within the IAM platform is essential to increase the number of fully automated fulfillment processes using models applied to transaction data. Most IAM platforms will push provisioning data to directory services. Pulling entitlement data from multiple directory services to apply new models based on entitlement usage patterns represents an IAM platform's future requirements. The IAM vendor platform in place today should accommodate a future road map with advanced analytics, and if it does not, then it's time to consider a new platform. The implementation and software acquisition costs will be offset with lower operating costs within two budget years. Fewer IAM staff to fulfill transactions are needed, with the remaining IAM staff learning more marketable skills (data analytics and model-driven security control design), resulting in sustained high performance.

It is now "cool" to be a part of an IAM transformation process. On the IAM transformation team, you will learn new, highly marketable skills and work with AI/ML models in big data applications, demonstrating more responsiveness to enterprise needs while managing the complexity of enterprise access management across tech stacks and cloud providers.

The New IAM- Model-Driven Security

The new model for IAM is built on a foundation of data science applied to a wide range of data sources, user attributes, models, access policies, and security controls across platforms. Data scientists will exploit increasing advanced analytics capabilities for enterprise benefits, including predictive models, model-driven security controls, and adaptive access policies based on behavioral attributes. Banks and credit card companies have used transactional behavior for operational fraud management for many decades, helping consumers and lowering online fraud. Incorporating transactional analysis into near real-time use of AI/ML models represents table stakes for maturing enterprises. IAM professionals need to master the fundamentals of data science to enhance IAM control design, improve user experience, and lower operating costs. IAM functions need data scientists for advanced analytics using models and improving data quality to help comply with privacy requirements.

IAM's foundation of data science is based on a simple construct we learned in our high school math class. We can represent data on a graph, specifically an X & Y axis chart. We can easily visualize a specific data point and then a cluster of data points. (See Figure 1) The cluster represents a pattern. The graph lets us visualize and measure the deviation from the pattern for a new data point (5.0, 9.0) in the diagram. Measuring the variation of a pattern is essential to applying the right automated workflow based on a threshold (a deviation score that is pre-determined). The threshold triggers an automatic action when the deviation reaches or surpasses the threshold.

Figure 1

Exemplar Graph of Data

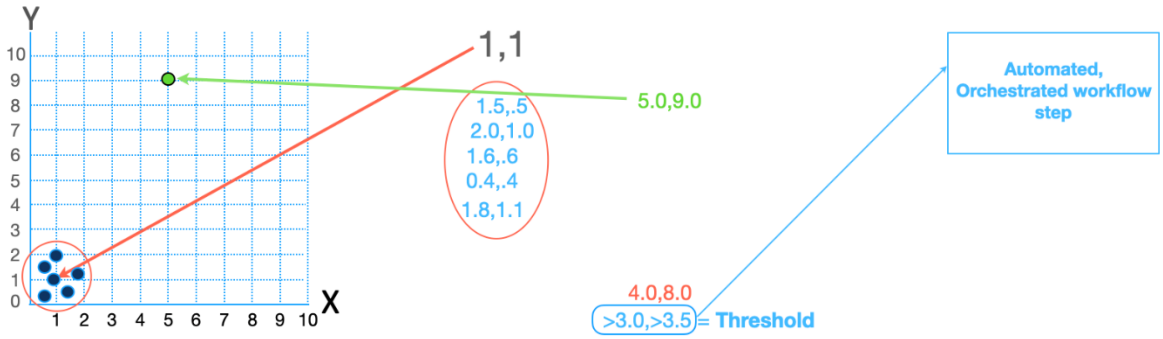

Figure 1

The key from a cyber perspective in IAM is to use this capability to define a normal behavior pattern for each respective user/consumer. Then, data points outside the pattern representing deviation at a threshold level will trigger automated action- perhaps revoking access to a system, as an example. Of course, this happens in milliseconds with computing capabilities available today, including streaming data architectures, data lakes, and cloud computing resources. It changes the game for IAM, where behavioral models can be applied to confirm deviation from a pattern and identify when credentials are compromised. Using AI/ML to determine patterns of threat actors has a significantly higher degree of difficulty since threat actors, by design, deviate from established patterns to break into a system. However, identifying normal behavior patterns for legitimate users is relatively easier due to the abundance of attribute information available regarding their behavior.

The FBI made public what the CIA figured out, that behavior doesn't lie. They attempted to identify microexpression changes during interrogations to determine when people were drifting from the truth. In the case of online usage patterns, it turns out that identifying deviation from patterns is straightforward. IAM has hundreds of viable use cases that can be enabled through AI/ML models applied to streaming data.

IAM has used secrets management (passwords) for authentication for many decades. Passwords have served the enterprise well for the past 60 years. Unfortunately, passwords are not the problem today. The number of digital assets requiring passwords is the problem. Users have too many passwords to remember, so they re-use passwords making it easier for cybercriminals to apply credential-stuffing techniques to obtain access to customer data. When all of us were taught IT in school or at work, we learned that authentication was an event with a beginning, end, and binary outcome. You were let into the system if you provided your user ID and password. If not, you had no access. The increasing risk of data compromise has forced enterprises to add authenticators to the authentication event with mixed results.

A better way to think about advanced authentication (that does not require passwords) is to think of authentication as a continuous process and no longer an event. The ongoing process can continuously feed transaction information based on multiple online attributes. The aggregate deviation score

(combined across attributes) can be fed in real-time to the system, consistently providing the deviation score compared to the established patterns of the attributes. Higher deviation triggers automated actions without the need for human decisions or actions. This happens in milliseconds while data scientists analyze transaction data for trend analysis.

Extending authentication into a continuous realm opens up opportunities to use online behavior to determine the right identity for every user/customer and whatever authenticators are used. Multi-factor authentication (MFA) is often associated with a user ID, Password, and additional factors like a one-time password shared through a text message. This conventional model for MFA has proven insufficient due to credential stuffing attacks and methods to spoof cell phone intercepting the one-time password. Multi-factor authentication is a useful descriptor that includes continuous behavioral authentication and password elimination. There are many choices in the market today for continuous authentication, and several that combine online fraud operational management orchestration to enhance ongoing risk management. All of them utilize data science to define behavioral patterns.

This same approach can be used for privileged user management (PAM). Historically, PAM solutions enabled the vaulting of a password the privileged user needed to produce to obtain the entitlement for a specific period. The new IAM approach will initiate a continuous monitoring process of the privileged user's online behaviors compared with the baseline pattern to ensure the privileged user is initiating actions and not a threat actor using the credentials of the privileged user. Once the deviation score moves above a pre-determined threshold, it triggers a revocation of privilege and the creation of a cybersecurity incident. This enhanced PAM model is impossible without applying data science and advanced analytics. It also does not require additive IAM staff to run this. Measuring user behavior at an enterprise scale is easily accomplished with the existing data management tools.

Behavioral models can be applied to provisioning access and access recertification resulting in fewer points of human approval, higher productivity, and lower operating cost. As an example, if a user pattern is established over a few weeks and that same user requests an additive entitlement capability (access to a specific data store), then the pattern can be compared to the pattern of a similar role, and if it is close enough (confirmed by the model); the provision request is fulfilled without human approval. The IAM team can monitor behavior deviation for specific users or many users and send an automated email or revoke access automatically based on the deviation score. Access recertification requests on an annual basis can be modified to eliminate the entitlement for users who never used it and approve it for those who used it within the acceptable range of behavior or risk score. This way, higher-level or riskier entitlement recertification requests can be routed to people for approval, while the vast majority are automated and not requiring human review.

A New Leadership Paradigm

IAM transformation requires a demonstrated commitment to the professional development of IAM resources. Leaders should ensure that all employees understand they will be asked to identify the skills they wish to develop/master while documenting them in a development plan. This practice reinforces the awareness that cybersecurity constantly evolves as threat actors adjust their tactics. Therefore, cybersecurity professionals should embrace the opportunity to learn and master different skills. The development plan offers employees a more explicit definition of the marketable skills they seek. At the

same time, the enterprise can share a commitment to pursuing the selected skills by identifying development activities to be incorporated into the work environment for each employee. Here is an example of a development plan:

Figure 2

Professional Development Plan

#	**Desired skill**	**Development activities**
1	Demonstrate the ability to apply an AI/ML model to an existing process resulting in a model-driven security control	
		1) Read selected material on cluster analysis
		2) Select a potential use case for applying a model to eliminate the need for human decisions/workflow
		3) Design the model and have an expert review and critique the model
		4) Apply the model and measure the results with a KPI
		5) Tune and adjust the model based on the KPI results

Employees are more likely to fulfill their development plans if they choose the desired skill to pursue through development activities. Their leader can help them by recommending development activities and adjusting to their role to enable them to learn how to apply the selected skill. Writing the skill down in the development plan helps employees focus on pursuing the specific skill rather than attempting to learn a set of skills broadly. A specific example is if an employee wants to be a world-class presenter. World-class level presenters have to master several skills to achieve that outcome, including presentation outline, design, and creating compelling content, in addition to mastering presentation delivery techniques. Leaders should encourage employees to be specific in defining the skill they wish to invest in at a granular level to make it easier to identify development activities. Gaining depth in development activities (10-12) is more effective in improving marketability. Marketability is simply the opportunity to have professional choices. More choices are better than fewer choices. Going deep into the number of development activities improves the options for an employee.

Highly effective leaders that demonstrate proficiency in transformative leadership skills are, more often than not, effective educators. The more significant the scope of the transformational initiative, the more emphasis on professional development is required. The leadership commitment to development heavily influences how employees respond to a transformational program. A strong demonstrated commitment to education by a leader will become an enabler for employees uncertain about their future and ambivalent about their careers in cybersecurity. Leaders acting as educators will deliver positive results in talent acquisition and development, helping IAM professionals better prepare for their future with the necessary expertise in AI/MI models applied to IAM functions.

Conclusion

Enterprise-level IAM built on a foundation of data science is what makes IAM cool today. The information shared in this paper is based on applied practices in several enterprises, so it is not theoretical. Top technical talent is attracted to the extensive use of AI/ML capabilities displacing older labor-intensive processes. IAM processes become more efficient and effective at a lower operating cost for the enterprise. The operational cost savings can be reinvested in more advanced use cases based on model-driven security control design and deployment. Leaders that are effective as educators contribute to the professional development of employees (talent development) while making it easier to attract new employees (talent acquisition) as a result of their commitment to employee development. Data science and the maturation of AI/ML models applied to streaming data architectures have opened up a tremendous opportunity for all IAM professionals to learn highly marketable skill sets based on data science fundamentals applied to cybersecurity. The enterprise benefits from more effective IAM controls deployed at a lower cost that greatly improves the user experience across the enterprise. The key for the CISO is to recognize the need to develop data science skills organically and through partnering with colleagues and acquiring talent. Demonstrating the ability to be an effective educator is one of the keys to leading a successful IAM transformation for any enterprise. In other words, make IAM cool!

About the Author

Jim Routh is currently serving on several boards and advising several companies. He is the former Board Chair of the Health Information Sharing & Analysis Center (H-ISAC), where he served for five years, and a former Board member of the Financial Services Information Sharing & Analysis Center (FS-ISAC). Jim is a former CSO/CISO for American Express, DTCC, KPMG, Aetna, CVS, and MassMutual. Jim brings a vast business and technology background to the boards he serves and is considered a digital and cyber security industry expert and thought leader. Jim is an ICIT Fellow and an Adjunct Faculty member, where he teaches cybersecurity for the NYU Tandon School of Engineering.

About the Organization

The Institute for Critical Infrastructure Technology (ICIT) is the nation's leading cybersecurity think tank providing **objective, nonpartisan research, advisory,** and **education** to **legislative**, commercial, and public-sector cybersecurity stakeholders.

ICIT understands that only through generative and focused collaboration will cybersecurity and national security communities make the quantum leaps necessary to defend against today's hyper-evolving

adversaries. In response, we facilitate a robust platform of programs, knowledge sharing, cutting-edge research, and publications that support the exchange of ideas and provide a forum for cybersecurity leaders to engage in the meaningful discourse needed to effectively support and protect our nation's critical infrastructures.