SANS

Analyst Program

White Paper

Identity Governance and Administration Powered by Risk Context: Taking Access Control to the Next Level

Written by **AJ Yawn with Terry Hicks**

October 2023

clarity

Introduction

Identity governance and administration (IGA)—ensuring that the right people have the right access to systems and data at the right time for the right reasons—has emerged as a business-critical concern for enterprises and their security organizations worldwide. IGA is key to protecting sensitive systems and data, and that makes it the first line of defense against acutely damaging data breaches, malware attacks, identity theft, and many other kinds of security failures. A survey by the nonprofit Identity Defined Security Alliance (IDSA) showed an astonishing 90% of the respondents' enterprises hit by identity-based breaches in 2023. Phishing attacks were, by far, the leading cause, followed by brute-force attacks and password social engineering.[1]

Identity and access failures represent an urgent problem for enterprises, one that's growing in terms of both frequency and severity of impact. The list of high-profile enterprises that have been hit and hit hard by identity-based data breaches and other attacks is seemingly endless, and it just keeps getting longer and longer. It includes entities as diverse and wide-ranging as Adobe and Sony Pictures Entertainment, LinkedIn and Target, the Canada Revenue Agency and the U.S. Department of Defense. The damage these incidents cause is far-reaching and highly impactful. It includes legal liability—for both enterprises and their key decision makers—regulatory compliance failure, reputational damage, loss of trust, and much more. That makes identity and access an issue that enterprises have no choice but to take *very* seriously. And it's an enterprise-wide issue—not just a "security problem."

IGA is an enterprise-wide issue, not just a "security problem."

Let's take a hypothetical but entirely realistic look at one common way an identity-based attack can occur and the damage it can do. Alexei S., a highly skilled hacker working from Romania, begins probing the networks of a New York–based multinational investment bank, looking for weak spots in its defenses. He uses advanced evasion techniques and encrypts his traffic so it blends in with normal activity. In addition, he uses a low-and-slow scanning approach, spreading his probes over a long period so he won't trigger any rate-based intrusion detection systems (IDSs) the bank may have installed. His attention to detail pays off. He identifies an employee, John A., who has unusually high-level access privileges for someone in a nontechnical role and thinks a lack of sophistication about information security might make him an attractive target for a "spearphishing" attack. He's right. John clicks on a link in a spoofed email about an "urgent security alert"—a link that downloads a trojan, a virus that disguises itself as legitimate code. Alexei suddenly has access to an astonishing range of sensitive data, from personnel files to intellectual property (IP) to company financials to client accounts. Alexei begins siphoning funds from client accounts, avoiding detection by taking only comparatively small amounts from any one account, and working only after U.S. business hours. By the time a routine audit identifies discrepancies in the accounts

[1] "2023 Trends in Securing Digital Identities," www.idsalliance.org/white-paper/2023-trends-in-securing-digital-identities

and an investigation is begun, he's stolen more than $10 million, transferring it through a series of offshore banks that make the money trail effectively untraceable. Alexei is long gone too, and so is his trojan, which he programmed to delete itself after 90 days. Now the bank is facing a Securities and Exchange Commission investigation and a slew of class-action lawsuits by both its clients and its shareholders.

IGA is central to preventing security failures like these, and to managing identity and access efficiently and effectively enterprise wide. But IGA—and identity and access control in general—are becoming more challenging, more difficult, and more burdensome all the time. There are a lot of reasons for this worrying trend. They include the growing complexity of enterprises' IT environments, changes in the way enterprises operate and their employees work, massive increases in the number and type of identities that need to be managed, and increasingly rigorous regulatory compliance requirements, to name just a few of the most important.

These are enterprise-wide problems, and they impact a wide array of enterprise stakeholders. Security, IT, and business managers alike are struggling with an identity and access environment that's simply become too complex, too big, and too dangerous to be addressed manually.

That's why a new class of IGA technologies, powered by risk context and using advanced machine learning (ML) capabilities, has emerged to automate these processes—completely disrupting the IGA market in the process. These advanced tools take into account complex contextual factors such as who wants access to what systems and data, when, where, why, and more. They use ML to make decisions based on those factors. It's increasingly clear that it's the only way to ensure IGA is truly effective in today's fast-changing business and IT environment.

Identity and Access Failures in the Real World

Security failures caused by poor identity and access controls aren't purely hypothetical, of course. Here are just a few real-world examples of recent highly damaging cases:

- **Equifax—A 2017 data breach at the credit reporting agency—caused by the failure to patch a known security flaw—compromised the personal data of more than 150 million people in the U.S. and elsewhere. The cost to the company, in regulatory fines, lawsuit settlements, legal costs, and more, is estimated at more than $2 billion to date.**
- **Marriott International—An email spoofing attack caused the theft of the personal data of more than 500 million of the hotel chain's guests in 2020. The cost to Marriott: a $124 million fine from the U.S. Federal Trade Commission.**
- **British Airways—The personal information of more than 400,00 British Airways customers was exposed in 2018, in what is believed to have been a supply chain attack via a third-party payment site. The U.K. Information Commissioner's Office fined the airline a record £183 (approximately $230 million), which was later reduced to a still-significant £20 million (approximately $25 million) due to the economic damage caused by the COVID-19 pandemic.**
- **Memorial Healthcare Systems—In 2017, a former employee's login credentials were used to access the electronic protected health information (ePHI) of this Florida-based hospital operator. Investigators found widespread, longtime violations of the U.S. Health Information Portability and Accountability Act (HIPAA), with more than 100,000 individuals' ePHI stolen and, in some cases, sold. The failure cost Memorial a $5.5 million settlement with the U.S. Department of Health and Human Services.**

What Is IGA Powered by Risk Context?

This concept dramatically extends traditional IGA, by taking into account a complex array of contextual risk factors to make fine-grained decisions about who should be allowed access to what systems and data, for what reasons, and under what circumstances. It does this using advanced machine learning (ML) capabilities, both in making access decisions and, where appropriate, escalating and disclosing anomalies and other potential security problems.

The Trouble with Traditional IGA

Before we discuss the new IGA technologies that promise to take identity and access control to the next level, let's take a closer look at some of the reasons *traditional* IGA tools and processes simply aren't working anymore:

- **Complex, heterogeneous IT environments—**There's nothing unusual today about enterprises mixing cloud computing—including multicloud services from multiple providers—with on-premises implementations. Bring-your-own-device (BYOD) policies mean end users are likely working on Windows, MacOS, or Linux computers and connecting with a broad range of mobile devices. Open source server software like Apache is everywhere. Many enterprises are still using legacy systems, some of them developed in-house, that have specialized requirements. Data may be stored and managed locally or in the cloud. The autonomous sensors and other devices that make up the Internet of Things (IoT) are increasingly making operational decisions without human intervention. (It's worth noting that one key driver of all this complexity is merger-and-acquisition activity, which typically results in a patchwork quilt of systems and processes.) These all represent potential target vectors for hackers, and they all require highly specialized identity and access controls.
- **Changing business and operational models—**Most enterprises today are highly decentralized, with employees, functions, and processes widely distributed, both geographically and organizationally. One example: There was already a worldwide shift to remote/hybrid work models, and the COVID-19 pandemic made it clear that it's here to stay. Supply chains are more complex and, as the pandemic also showed, more fragile than ever. Enterprises frequently operate in large-scale industry ecosystems that may involve sharing systems and data with hundreds of business partners, and even with direct competitors. And all this complexity and change means that managing access privileges is more important and more critical than ever.
- **Increasingly rigorous regulatory compliance requirements—**A broad range of regulatory compliance frameworks, court decisions, and industry standards bodies are placing intense pressure on enterprises not only to improve the security of their systems, their data, and their business processes, but also to demonstrate that they're doing so. The SEC, for example, recently released a new set of cybersecurity rules that require publicly traded companies in the U.S. to disclose cybersecurity incidents within four days of having been determined to be "materially relevant" (that is, having an impact on investors' decision making). Another piece of U.S. legislation, the Sarbanes-Oxley Act, imposes rigorous financial transparency standards that can only be met with effective cybersecurity controls. And the European Union General Data Protection Regulation (GDPR)—the toughest set of privacy rules in the world—places serious restrictions on how personal and enterprise data can be handled, and even where it can be stored.

The result of all these factors, and many more, is that identities are literally out of control, in terms of both their number and their complexity. It's now entirely commonplace for a single end user to have multiple identities—sometimes dozens or even hundreds of them—with multiple levels of access and variable access privileges. Even machines, like robots on assembly lines, have their own identities.

This is a problem that extends far beyond the security organization. Security practitioners, risk owners, and business managers are all struggling to manage. Business managers are especially stressed, because they don't have the time, the skills, or—crucially—the contextual information to make effective decisions about identities and access privileges. And yet, they're likely to be held responsible for identity-based security failures and the resulting damage. This is clearly an unsustainable situation.

> **The complexity and scale of enterprise identities have become so great that they can no longer be managed manually. Automated IGA, powered by a fine-grained understanding of risk context, is now essential.**

The Next Step Forward in Identity and Access Control: IGA Powered by Risk Context

IGA is, of course, fundamentally concerned with managing risk—specifically, the risk that comes from allowing a specific individual or entity access to a specific system, process, or dataset. It's important to recognize that there's no such thing as a zero-risk environment. It's both unrealistic and undesirable to try to eliminate risk entirely. Enterprises always must accept a certain degree of risk because they can't operate, remain competitive, or even survive otherwise. Instead of trying to get rid of risk, they—and their security organizations and other stakeholders—should work to achieve an appropriate and acceptable balance between risk and reward.

When we talk about IGA powered by risk context, it's important to understand two closely related but discrete concepts:

- **Inherent risk is the risk.** The potential for damaging impact—that is inevitable in a given environment. (It might be useful to think of it as "risk in a vacuum.") In cybersecurity terms, this could mean anything from a lack of security awareness to unpatched applications to the physical failure of servers. Inherent risk is comparatively easy to recognize, predict, and address, for example, with employee training, regular software updates, and equipment checks. It's the primary focus of traditional IGA technologies, and it's that focus that represents the primary weakness of those technologies.
- **Contextual risk is influenced by the specific context in which an enterprise operates.** Contextual factors—everything from the social, political, and economic environment to emerging cyber threats to the physical locations of employees and the times they're likely to be working—are intrinsically complex and constantly changing, and they're far more challenging to deal with. But recognizing and mitigating contextual risk is critical to protecting the enterprise today.

Just as enterprises must balance risk and reward, they also must balance inherent and contextual risk. (Many industry observers consider an appropriate balance to be approximately 60% inherent and 40% contextual.) As we've already noted, getting contextual risk *right* is the more difficult part. Contextual risk includes an extraordinary range of factors. Here are some of the most important contextual factors to be incorporated in an access determination:

- **Type of systems and data—**Different systems—say, a financial application and a marketing database—have different risk profiles. The type of data, for example whether it's sensitive personal information, IP, or financial records, also plays a role.
- **Individual roles—**The function of the person requesting access is crucial. An HR representative likely needs access to personnel files, but not necessarily to financial records.
- **Time of access—**The timing of an access request (for example, late at night) can be a red flag.
- **Location of access—**Access requests from locations not usually associated with a given user could be considered risky.
- **Purpose of access—**Understanding *why* the user wants access—whether it's for routine work, a specific project, or some unspecified reasons—can help evaluate the associated risk.
- **Previous behavior—**Past behavior can be an excellent predictor of the legitimacy of a request, and anomalies in behavior patterns can trigger request escalation and other additional security measures.
- **Current security posture—**The overall security of the enterprise's systems, network, and data at the time of a request also can influence an access decision. When a known wave of cyberattacks is ongoing, for example, any access request might be treated with greater-than-usual caution.

Advanced IGA technologies, powered by risk context and automated using ML capabilities, can make decisions based on these complex and multifaceted factors more rapidly, more efficiently, and more accurately than human beings. They also can do it at a scale that's impossible for human beings. For example, the technology can identify subtle anomalies that would likely go unnoticed by a human being, particularly one, like a business-level manager, who lacks specialized security skills.

Here's an example of how this could work in real life. Sarah M. is a security manager at a healthcare provider that operates more than a dozen hospitals and clinics in California, Nevada, and Arizona. She has been tasked with ensuring the security of the company's highly sensitive patient information, which is, of course, subject to stringent HIPAA regulations. The company's senior management is extremely worried about this issue because several similar healthcare operators have been hit recently by serious data breaches—one of them a crippling ransomware attack. Sarah knows the identity and access controls in place are inadequate, but she also knows improving them won't be a simple undertaking. Thousands of users, ranging from physicians conducting telehealth visits to developers upgrading mission-critical applications, need access to different systems using different devices from different locations. Because the company operates in multiple states, it will likely have to address varying regulatory compliance requirements. And everything needs to function seamlessly, even when an individual requesting access isn't very technically knowledgeable. Sarah decides to implement an automated IGA application powered by risk context. (This isn't the only step she takes, of course. She also institutes regular access reviews and implements enterprise-wide multifactor authentication.) The new IGA system balances inherent risk against contextual risk, then uses ML to automate the process of determining access. It doesn't take long for the new application to detect questionable activity. A data analyst has logged on to search for a specific patient file. As an analyst, he has privileges for the records in question, but it's highly unusual to request just one. ML flags the inquiry as questionable, blocks it, and escalates it to a security manager. It's a good thing. It turns out the analyst is involved in a bitter divorce and custody case and was looking for information about his estranged wife. If that data had been leaked, it would have been a serious violation of HIPAA and other privacy laws and would have exposed the company to serious legal liability. The analyst is, of course, immediately terminated—but not before his access to all systems and data is cut off. Sarah can report to her management, to the company's governance committee, and to the relevant regulatory authorities that a data breach has been successfully averted without any data being compromised.

Identity and access failures like this hypothetical scenario are all too common, and they can only be prevented using advanced tools that can input contextual factors and use them to detect anomalies—and take action without human intervention.

Identity governance and administration isn't just about who should be allowed access to the enterprise's systems. It's also about who wants access to what systems and data, when, where, and—most importantly—why.

The Benefits of IGA Powered by Risk Context

The new IGA tools coming on the market, powered by comprehensive risk context and ML capabilities, have the potential to dramatically improve the security of highly sensitive enterprise systems, data, and processes. That sharply reduces the risks of identity theft, data breaches, ransomware, and other cyberattacks, inadvertent loss of personal data or intellectual property (IP), and many other identity-related security failures. That, in turn, reduces the attendant risks of regulatory compliance failure, legal liability, reputational damage, and loss of trust by customers, partners, and the public.

It's almost impossible to overstate the importance of improving identity and access controls. They're key to protecting the enterprise—and also many stakeholders inside and outside the enterprise. As we've already noted, regulators and lawmakers are more and more willing to hold board members and other risk decision makers responsible for failures of security and risk management. The courts, too, are becoming more and more involved. It's also becoming clear that existing IGA and other identity and access controls are inadequate to protect them from these serious risks.

The benefits of this new approach to IGA aren't limited to security improvements, however. IGA powered by risk context holds the promise to reduce the burden of access control for business managers, who—as we've seen—don't have the contextual knowledge or the subject-matter expertise to make complex decisions about privileges. These advanced tools also will make it easier for them to lower their risk scores, for example, by cleaning up orphaned accounts and removing other types of improper access.

IGA powered by risk context also can make life easier for security organizations by automating previously burdensome processes—especially by reducing the number of ad hoc privilege reviews that must be conducted manually—and freeing up security personnel to perform higher-level value-added functions.

The Bottom Line

Traditional IGA is no longer adequate to meet the identity and access needs of the modern enterprise. A complex, fast-changing IT and business environment is driving an unmanageable proliferation of identities—for individuals, for organizations, even for devices—with equally unmanageable privilege requirements. Trying to manage all these identities, competently and in a timely manner, is becoming impossible for traditional IGA tools, and is placing excessive and unreasonable burdens on business managers, risk owners, and security organizations. The answer to this urgent set of problems is IGA powered by risk context and enabled by ML. This new set of technologies represents the future of identity and access control—and a major step forward for enterprise security.

Who Can Benefit from IGA Powered by Risk Context?

Security failures caused by poor identity and access controls aren't purely hypothetical, of course. Here are just a few real-world examples of recent highly damaging cases:

- **The CISO, and everybody else in the security organization—**Security practitioners at all levels have the assurance that systems and data are protected. This is especially important because it's the security organization that's most likely to be blamed for an identity-based security failure, and the consequences can be very serious. One example: The CISO of Uber was convicted of criminal charges for his role in covering up a data breach that impacted more than 50 million customers and drivers. He was fined $50,000 and ordered to do 200 hours of community service—and, as the judge in the case made clear, he only narrowly escaped prison time. The benefits of advanced IGA practices don't end there. They also can free up front-line security practitioners to take on more rewarding value-added functions.
- **The board of directors—**The benefits of risk-context-powered IGA reach all the way to the very top of the enterprise. The new SEC rules make it clear that board members can and will be held *personally* responsible for cybersecurity incidents, for the failure to disclose them in a timely manner, and for not having appropriate governance standards in place. A long series of legal decisions shows that courts are increasingly prepared to impose serious civil and potentially even criminal penalties. IGA enhancements will offer them important protections.
- **The CEO, senior executives, and line-of-business leaders—**These are the people responsible for the enterprise's day-to-day operations—and the ones who'll be held responsible for the operational impacts of identity-based security failures.
- **Business managers—**These front-line managers are all too often tasked with making decisions about access privileges—frequently on an ad hoc basis—that they have neither the subject-matter expertise nor the contextual information to make. This places an excessive management burden on them and places the enterprise at unnecessary risk.
- **Auditors (internal and external)—**The more contextual information auditors have about the enterprise's identity and access controls, the more credibility their security audits will have and the more confidence investors will have in the enterprise.
- **The chief risk officer (CRO) and chief compliance officer (CCO)—**These two closely related roles have an obvious interest in identity and access because they're critical components of risk management, regulatory compliance, and corporate governance.
- **The chief information officer (CIO) and the chief technology officer (CTO)—**These high-level technology decision makers have enough to do just managing constantly evolving IT and business environments. They don't have the bandwidth to be making detailed decisions about access privileges.

Sponsor

SANS would like to thank this paper's sponsor:

clarity