Book Review & Analysis
Identity Attack Vectors, by Morey J. Haber & Darran Rolls
Reviewed by Tom Deaderick

Objective(s)

- Summarize the book, extracting the most useful elements for others
- Relate the book's methods to our processes and challenges for discussion and implementation.

Legend

Excerpts from the book are in black with page numbers.

Commentary is in blue.

This book also includes descriptions of features which are potentially helpful in evaluation, or selection of an IGA system or useful in prioritizing features of an upgrade. **These are orange and bolded.**

Table of contents

Foreword

"Consider this for a moment - no one at Amazon has ever physically seen me, verified my account information, or validated my signature." (p. xvii)

Technologies outside the office continually adjust expectations of HCA patients, business partners and employees. What was leading-edge yesterday becomes an accepted norm today. This sentence has ramifications in more than one direction.

Assigning a system identity to a person (e.g. "onboarding") previously required the person to present physical identity credentials in person. NIST 800-63A allows people to claim identities through remote procedures, including the scanning of State-issued credentials. As these requirements relax, the need to guard against attackers that seek to acquire privileged identities increases.

The second ramification of the sentence is that while Amazon has never physically seen most of their millions of customers, it has a wealth of data describing each user's preferences, behaviors and buying preferences. These attributes are more important to Amazon, than matching the account to a user's physical appearance. The user's digital identity is more important to Amazon than their physical identity.

> "So, after those considerations, one must wonder why, as a security industry, we spend so much effort on protecting networks and servers and comparatively so little attention on understanding and protecting our user's identities. For decades, we have poured countless billions of dollars into perimeter protection, and yet, year over year, the number and severity of breaches continues to rise. Each year brings new attack vectors that we need to guard against, but yet, the goal of the attacks is usually the same; obtain account credentials, elevate the privileges of those accounts, and steal as much information as possible." (p. xviii)
>
> "With the network perimeter continuing its dissolution, and an almost obsessive move to cloud-based services, **the user's identity is one of the last bastions of control still in the hands of enterprises**." (p. xviii)

This is a strategically clarifying concept, yet as organizations recognize the need to invest in Identity, another factor bears consideration.

> "Unlike many other security projects, an Identity Management project typically breaks down almost every technology/business barrier within an organization. One of the biggest challenges many Identity Management projects face is the integration of business applications with user directories that are deep in the infrastructure, frequently trying to connect modern UI services with outdated legacy directories." (p. xix)

Chapter 1: The Three Pillars of CyberSecurity

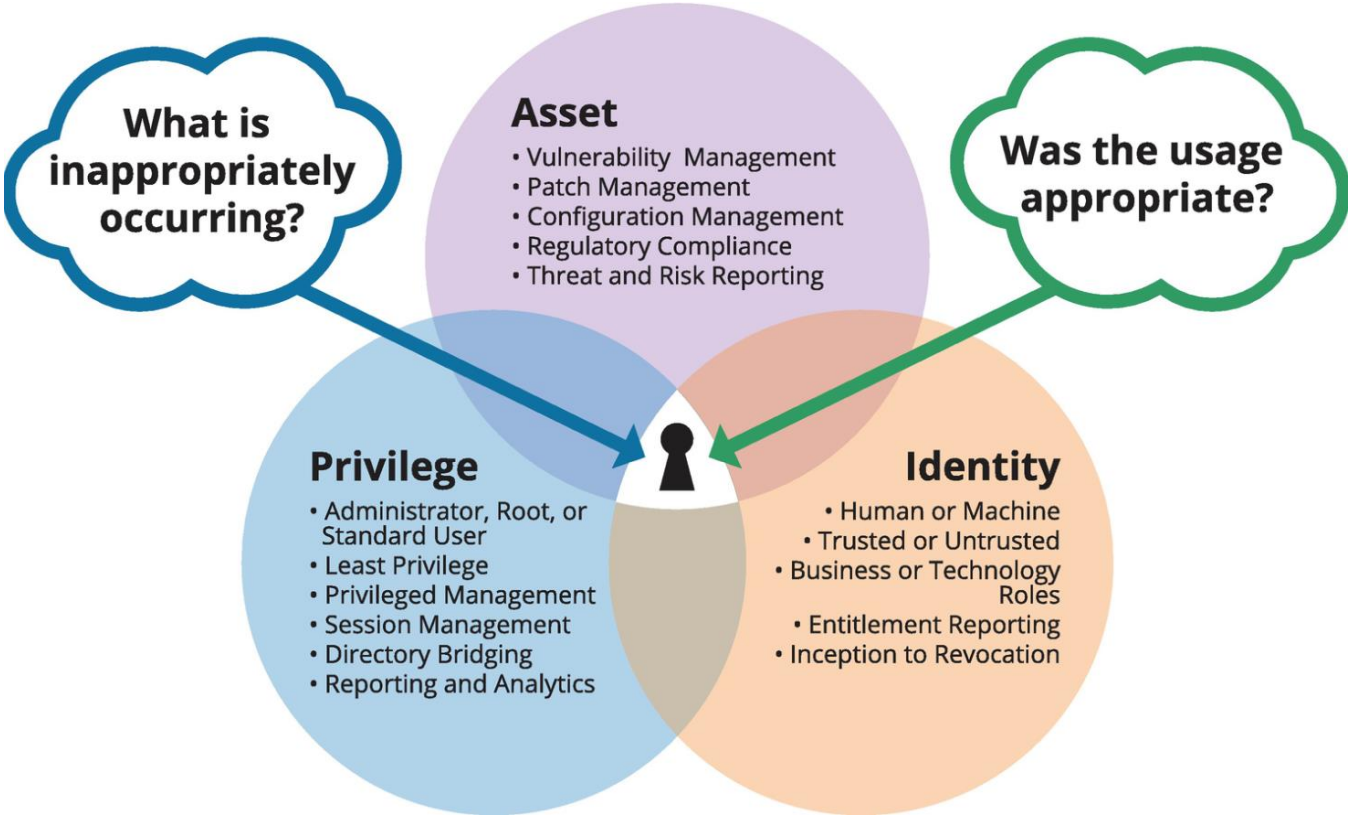

The book defines Three Pillars of CyberSecurity, which is more access-oriented than the more familiar Three Pillars of Information Security below.

Chapter 2: A Nuance on Lateral Movement

> "To a threat actor, lateral movement means all the difference between compromising a single resource and potentially navigating throughout an organization to establish a persistent presence." (p. 7)

> "Lateral movement, by the most traditional definition, is the ability to pivot from one resource to another and to navigate among other resources in any environment." (p. 7)

Controlling lateral movement is accomplished by limiting privileges of all accounts to only those essential for users to complete assigned tasks, a principle known as "Least Privilege".

> "Zero trust is a security model based on the principle of maintaining strict access controls and not trusting anyone, anywhere, at any time, even those already inside the network perimeter, by default." (p. 9)

> "**Just-in-time privileged access is a strategy that aligns real-time requests for usage of privileged accounts directly with entitlements, workflow, and appropriate access policies.**" (p. 9)

Chapter 3: The Five A's of Enterprise IAM

"One of the greatest challenges all security frameworks face is their complexity." (p. 11)

"The Five A's cover Authentication, Authorization, Administration, Audit and Analytics." (p. 11)

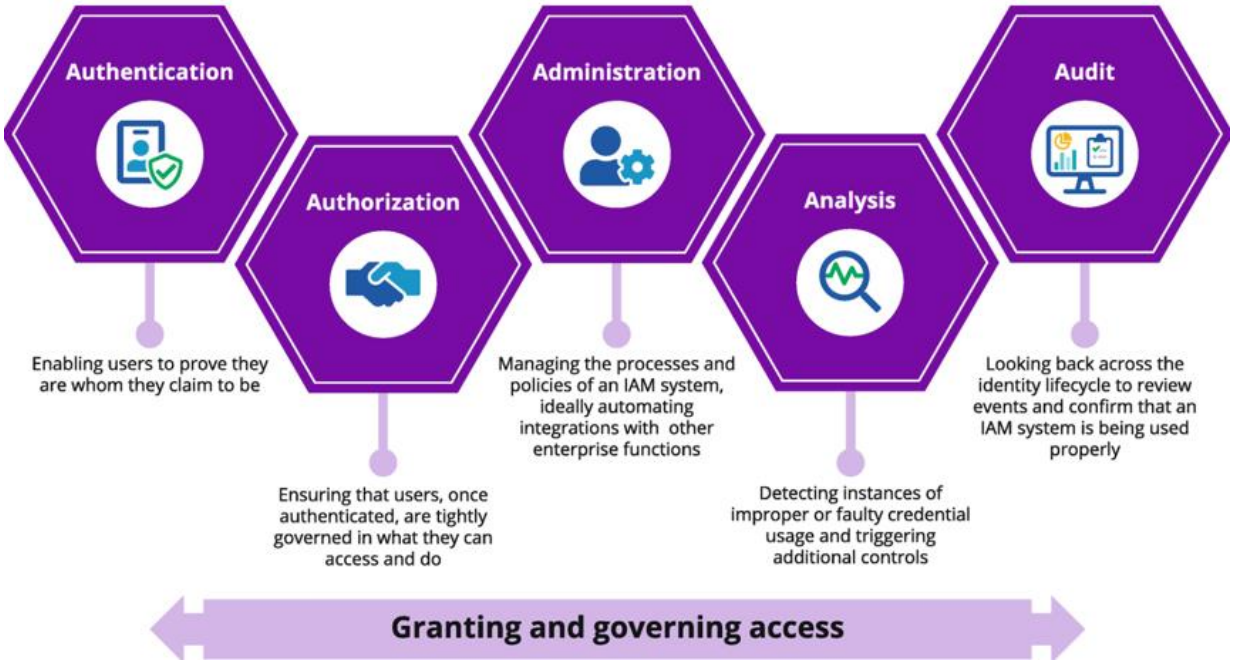

"**Authentication** is a login (username) in addition to some form of secret, historically a password, to establish proof or trust in an identity." (p. 12)

"You do not visually identify the privileges with the account simply by looking at the account or username." (p. 13)

This refers to avoiding the use of initials in usernames.

"So, in simple terms, **authentication** is nothing more than proving your identity or your ownership of a given account." (p. 13)

"**Authorization** is the right to perform a function based on your authentication." (p. 13)

"When like privileges are grouped together, they create the foundation of a role." (p. 13)

Administration in this context refers specifically to "*the administration of authentication, authorization, and audit controls*. Administration here means *configuration management and governance controls over any changes made to that authentication, authorization and audit*." (p. 14)

"**IAM administration is a big part of the Identity Governance remit. Add to it Audit and Analytics, and you have the product scope of most enterprise class Identity Governance solutions.**" (p. 15)

"For some, audit means delivering a user access certification program. For others, it's defining and implementing preventative and detective policy such as Separation of Duty (SoD). For all, it is being able to prove that comprehensive administration processes are in place and are being adhered to." (p. 15)

"**Advanced identity analytics enables a more informed and predictive approach to governance. Using Machine Learning (ML) and Artificial Intelligence (AI) techniques, identity analytics tools can provide important peer-group analysis information that helps to extend identity audit and administration functions and make them more dynamic and responsive.** For example, if an analytics engine discovers suspicious, inappropriate or unusual access, it can prompt administrators to review that access to ensure that the correct configuration has been implemented. **Analytics can provide auto-generated insights and recommendations that allow the line of business to make more informed access decisions that enhance operational security and ensure compliance.**" (p.16)

Chapter 4: Understanding Enterprise Identity

"Identity is typically a one-to-one relationship between a human being (a carbon-based life form) and their digital presence. Their digital presence, however, can have multiple accounts, multiple credentials, and an infinite number of entitlements in its electronic format." (p.17)

"A persona is a derivative of an Identity and refers to a special situation in which a person has multiple identities each of a different 'class'." (p.18)

"A physical persona refers to when a logical persona is reflected into the physical world. In the real world, there are sometimes physical traits that are connected to your persona. For example, these could be the uniforms we wear and the types and colors of badges we use for authorization." (p.19)

"An account is an electronic representation of an identity and can have a one-to-many relationship with the identity. One identity can, therefore, have multiple accounts. These accounts reference a set of permissions and privileges needed for an application or asset to connect or operate within the confines of a resource." (p. 22)

"Technically, any account is simply a vehicle to authorize usage and control operational parameters." (p. 22)

"A credential is an account with an associated password, passcode, certificate, or other types of key. Credentials can have more than one security mechanism assigned to them - this is called dual or multifactor authentication." (p. 22)

"Hacking an account means the same thing as taking control of its credentials." (p. 23)

"At the heart of every identity are privileges." (p. 23)

Users
"Users can have multiple accounts, credentials, and even personas, but they have only one identity." (p. 24)

"The representation for an identity should be a designator with little or no value and should not be directly linked to any other form of identification." (p. 24)

Account ownership
"Every identity requires an owner. If you are a human being, you are your owner. When an identity is assigned to an application or machine, it gets an additional attribute or owner or owners." (p. 29)

"The reason service accounts are a special type of account is because they should not have the same characteristics as a person logging onto a system. They should not have interactive users interface privileges nor the capabilities to operate as a normal account or user. Depending on the operating system or infrastructure, this could encompass restricting everything from executing a batch process to not having a proper shell assigned to the account. (p. 35)

"To reduce the likelihood and impact of attacks that abuse or escalate privileges, we should always strive to restrict the assignment of privileges to the lowest common denominator for every type of account. This concept is called **least privilege**." (p. 31)

Cloud accounts
"There technically is no accepted definition and terminology to explain cloud accounts. Each cloud provider - Software as a Service (SaaS) vendor, Platform as a Service (PaaS) vendor, and Infrastructure as a Service (IaaS) vendor - uses a different definition to meet their business and technology approaches to solutions and management in the cloud." (p. 36)

With cloud systems, "you can be compromised in two ways, unlike with an on-premise solution: first from the back, everything owned by the cloud provider, including their accounts used to support and manage the service." (p. 36)

Biometrics
"Identities linked with biometrics warrant special consideration since you cannot change things like your fingerprints (unlike a password)." If an attack exposes biometric data, "the threat actor has a method to authenticate you that can never be changed." (p. 37)

Entitlements
"Entitlements are any technology implementation that controls access to something we care to manage. Entitlement is a category name used for something that grants, resolves, enforces, revokes, reconciles, and administers fine-grained access, privileges, access rights, permissions, or rules." (p. 37)

Complex entitlements
"A good example of a commonly implemented complex entitlement is an SAP R3 role. By its nature, it is a composition of other entitlements (other roles, Tcodes, AuthCodes, etc. - each possibly entitlements in their own right), yet it is assignable as a single unit of access, so it is itself an entitlement." (p. 38)

Controls
"A control is a clearly defined management oversight function that enables the tracking of adherence to a given security or audit policy." (p. 38)

Roles
"At the highest level of abstraction, we define a role as a collection of people, or a collection of access, defined and maintained for the purpose of improved manageability, enhanced controls, and the promotion of good governance." (p. 39)

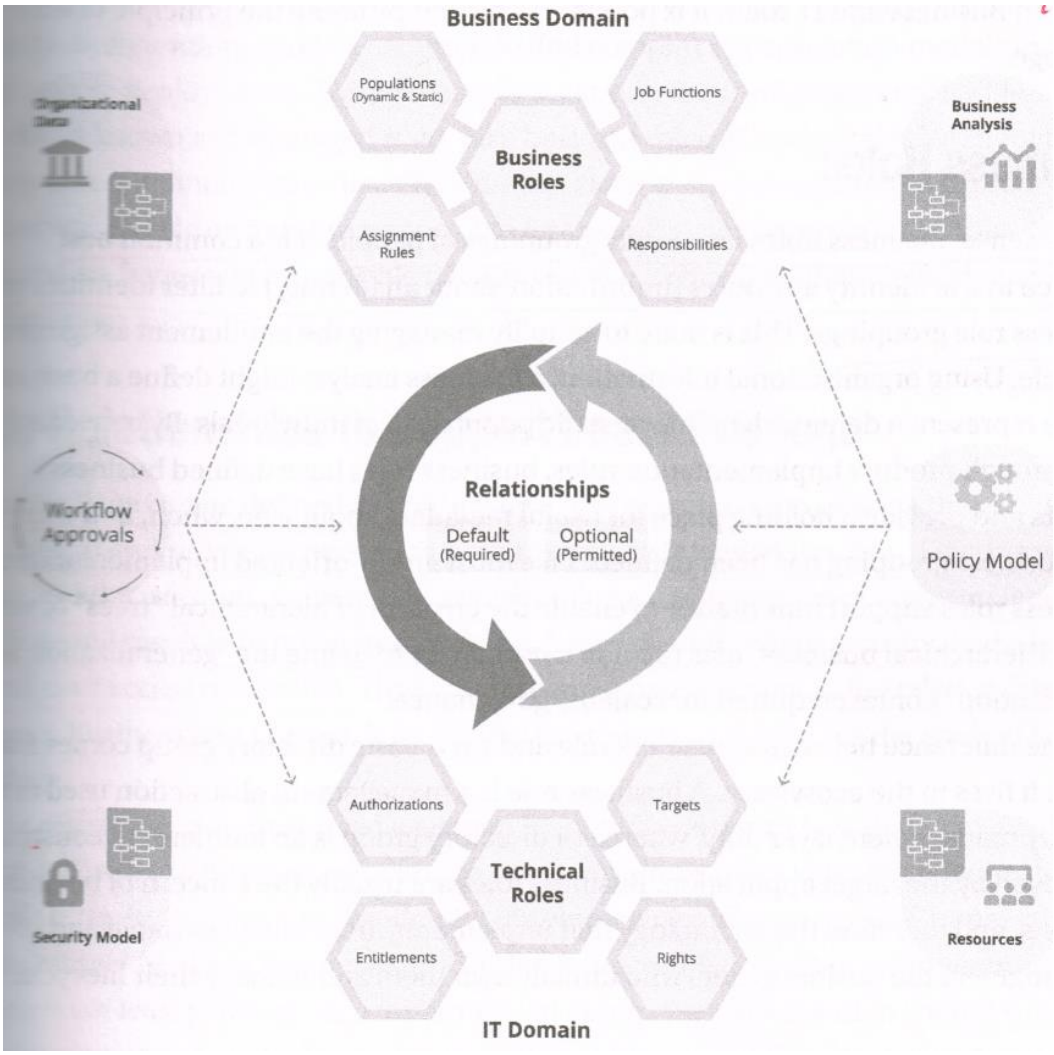

This illustration helps clarify some of the confusion in the various meanings of the word "role".

"Roles are sometimes used to collect like identities that perform similar functions and need the same level of access to technology assets. These roles are often referred to as Business roles. Roles are also sometimes used to collect related accounts and entitlements required to carry out a known set of actions. These are referred to as IT roles." (p. 40)

"Business and IT roles are connected together to form user assignment relationships. It's considered a general Identity Governance best practice to only connect IT roles to Business roles and use the Business role to identity relationship to complete the linkage back to the identity." (p. 40)

"In one sense, Business roles are simply groupings of people. It's a common best practice to use identity attributes (information about an identity) to filter identities into business role groupings." (p. 40)

Least Privilege
"Least privilege can be defined as only giving a user account or processing the entitlement and privileges needed to perform a given function. Steering your entire identity, account, and entitlement lifecycle management process toward least privilege will result in better system stability, better overall system security, and a lower overall user access risk profile." (p. 41)

Chapter 6: Identity Governance Defined

"Identity Governance is the technology and processes to ensure that people have appropriate access to applications and systems and that the organization always knows who has access to what, how that access can be used, and if that access conforms to policy." (p.45)

"We often use the term in Identity Governance 'JML' - this stands for Joiners, Movers and Leavers. This is a phrase taken from the realm of Human Resources and is used to capture the three major states that a typical user access process will go through." (p. 49)

Chapter 7: The Identity Governance Process

Authoritative Sources of Identity
"We use the term 'authoritative' to mean that these systems are the true source of user records for a given identity type or persona." (p. 52)

"Information about non-employees (contractors, business partners, and customers) tends to be stored in a mix of enterprise repositories like Microsoft Active Directory and custom database application systems." (p. 53)

"**A critical part of the IG process is consolidating these various authoritative sources to create a single view of all identity records. This single repository of user data does not replace the original authoritative sources; it simply creates a virtual consolidated view that we often refer to as a "governance system of record**". (p. 53)

"Be cautious of any identity and access administration system that divides its functionality based on 'where the app runs'. The end user does not care where the application 'runs', nor should you." (p. 54)

"**As a best practice, we highly recommend the vaulting and management of these (IGA) credentials in a PAM vault or password safe or API key management system that is external to the Identity Governance engine.**" (p.54)

Direct-API connectivity

"**Direct-API connectors provide connectivity between the IGA server and the target application by some form of API or remote read and write mechanism. In general, these connectors are provided by the IG vendor and should provide coverage for all of the major enterprise applications on-premises and in the cloud.**" (p.55)

Shared-repository connectivity and deferred access

"**Shared-repository connectors cover centralized systems like enterprise directories (.e.g. Active Directory), Single Sign-On systems (e.g. Okta), and all forms of externalized authorization.**" (p 55)

"In this model, a group membership is used to control functional access inside the application. The access control is therefore deferred from the application to groups and group membership with the centralized service. (p.56)

"The challenge then becomes sorting out which group membership (entitlement) belongs to which application. It's important to understand that to the directory, these are just accounts and groups, and there is nothing to relate them back to the actual applications." (p. 56)

Standards-based connectivity

"**Good examples of standards-based connectors are for systems like LDAP, JDBC, CSV, REST, and SCIM. Here the target system supports a standards-based connectivity API.**" (p. 56)

"**Another example of an important standards-based connector is the System for Cross-domain Identity Management (SCIM). In an ideal world, every application would have full support for SCIM.**" (p. 57)

Connector reconciliation and native change detection

"**The connector tier is responsible for reaching out into the infrastructure, making changes, and thus managing accounts associated with an identity or an owner. However, changes often happen locally to that application or infrastructure. Local admin actions can and do happen, and it's the job of the IG system to understand those changes and 'do the right thing'. This is the process of reconciliation and change detection.**" (p. 57)

"There are two primary methods of understanding that something has changed out in the real world. **Either the IG system is notified of the change by the target system or it does a delta change analysis based on its own cached (previous) version of the configuration - this is traditionally referred to as reconciliation or recon in the IG vendor space. It is important that your IG system and processes support both methods. If available, change detection from the managed system itself is preferable.** For example, most LDAP servers support a special attribute polling mechanism (in Microsoft Active Directory, this is referred to as 'USN-changed') that tells the reader that something has changed." (p.58)

"**Historically, reconciliation and change detection have been a significant area of product differentiation. Outside of processing cost of delta change analysis, some systems do not handle reconciliation well. When the reconciliation process finds delta records, these 'badly behaving implementations' have no choice but to treat the local change as an error and automatically change it back to the previous known value. This can cause unforeseen issues with product maintenance, upgrades, and even runtime security. Mature IG systems will allow for change triggers and control processes to be executed upon local change detection.**" (p. 58)

Correlation and orphan accounts

"The ongoing process of connecting people to the accounts and access is called correlation. In the ideal world, every account matches up perfectly with a human (identity), and you have 100% correlation (for the record, that's something we never see out of the gate). Account access that does not correlate to a known user is often referred to as an orphan account. Orphan accounts can be a significant security weakness. Post breach forensic analysis shows that the adversary creates and uses new accounts throughout the cyber killing chain. It is therefore essential for ongoing governance and security to instrument, and if at all possible, to automate, the detection and rapid resolution of orphan accounts." (p. 58)

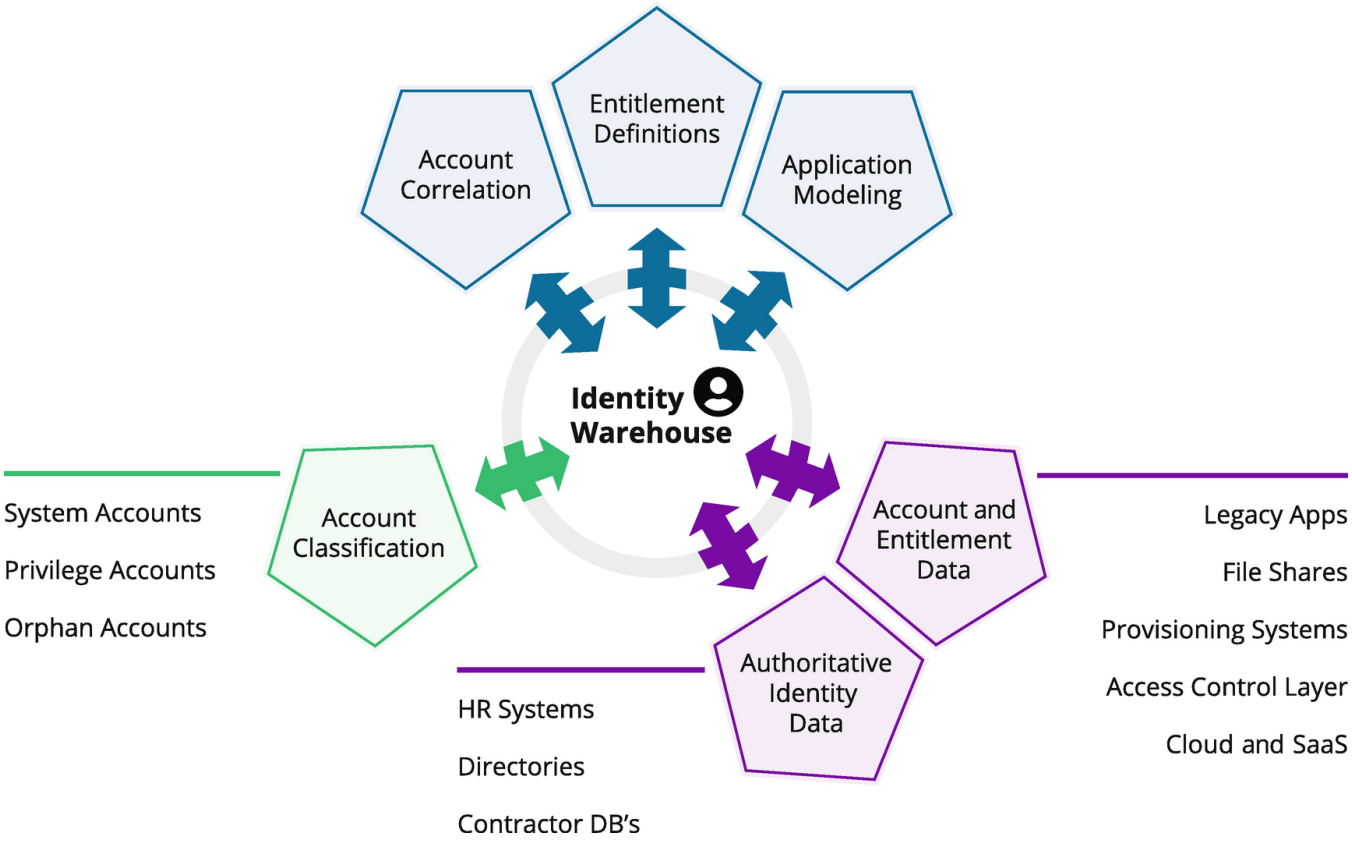

"Identities, accounts, and privileges continually flow into the system, and items that don't connect back to a human are flagged for the attention of application and system administrators." (p.59)

"In large ecosystems, there can be hundreds and potentially thousands of accounts that will not correlate without a deliberate and specific process of managed correlation. **An enterprise-grade IG solution will provide core product capabilities to help either manually or automatically resolve these issues. Manual correlation using graphical 'searching and connecting' will greatly help the admin establish and maintain links between known owners and orphan accounts. Automated matching algorithms can also help suggest relationships and potential connections.**" (p.59)

Building an entitlement catalog

"At the center of the Identity Governance (IG) process sits the entitlement catalog. An entitlement is the generic name given to a technical access control facility that we care to catalog and manage." (p. 60)

"We use the term entitlement as an abstraction for all things that provide access." (p. 60)

"**When selecting a commercial IG solution, we recommend looking for an entitlement catalog that delivers best practices around defining ownership, approval process, definitions, and classification capabilities. The entitlement catalog is the center of an**

IG solution, so look for a highly extensible metadata framework as part of the solution. This metadata will allow you to define custom attributes that represent your business's needs and satisfy the business requirements for regulatory compliance." (p. 60)

"A flexible entitlement catalog will greatly help drive the zero trust cause by allowing for the 'promotion' of basic identity attributes like location, or job code to be entitlements in the catalog." (p. 61)

"**Remember that, in an entitlement-driven control system, metadata is king; so, protect your metadata like it is the crown jewel. It literally is the data map for all identity-based entitlements within your entire managed environment.**" (p. 61)

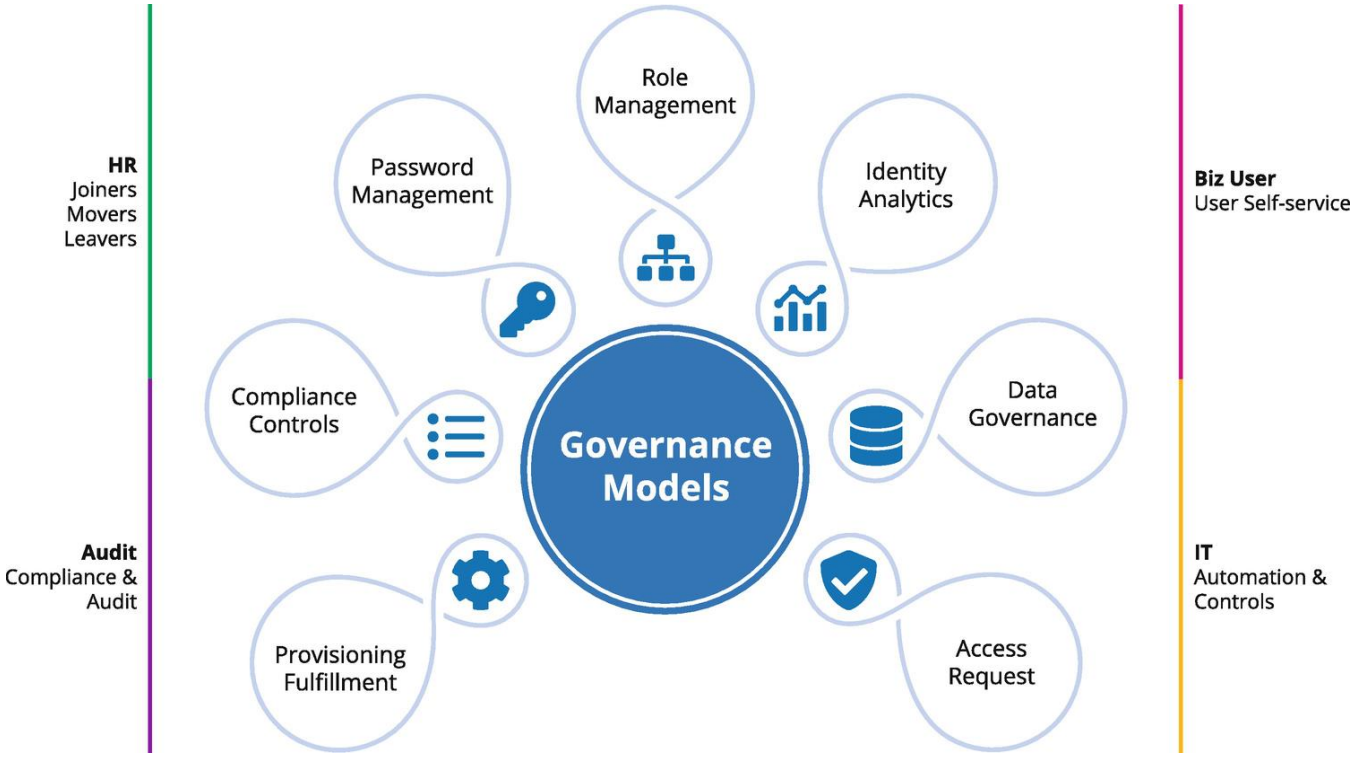

Lifecycle management

Sample LCM states

- Prehire
- Hired
- Terminated

"**IGA policy models are used to capture the desired state, known best practice configuration, and an inventory of controls and governance actions. These models are abstract representations of how accounts and privileges should be set, approved, audited, and used to some known state. Examples of IGA policy models include the entitlement catalog, provisioning schemas, approval and ownership records, audit requirements, role models, lifecycle triggers, and separation of duty rules.**" (p. 66)

"There's a phrase often used around successful IGA deployments - it goes something like 'let the models drive the process'." (p. 66)

Enterprise roles as a governance policy model
"Enterprise roles are a critically important model to get right. You don't have to use them to operate a governance-based lifecycle, but if you do use them and you get them right, you can vastly simplify the whole process." (p.66)

"Simply understand that a good role model will provide a place to define, verify, and reconcile access, a place to define the correct configuration or entitlement, a place to establish assignment approvals, and a place to track the ongoing state of access across potentially thousands of target applications, hundreds of thousands of users, and millions of entitlements." (p.66)

"Provisioning is the term long used in identity management to represent the overall process of delivering access to applications and data." (p. 67)

Should a replacement IGA system be considered, integration will obviously be a key criterion. The first, and perhaps most important system for the new IGA system to integrate successfully with, would be the existing system. Integration with the current system would allow a phased replacement rather than a difficult big bang implementation. This also enables access to the key features of the new IGA before the full system is deployed.

Account policies
"Policies are often required to manage accounts that have not been used for a specific period of time (usually referred to as dormant or stale accounts). Dormant accounts are a known common identity attack vector and are of specific concern to IT security." (p.72)

Entitlement policies
"The 'entitlement context' allows you to build policies that enforce specific assertions about who should and who shouldn't have access to the systems and data we are responsible for. For example, an entitlement policy might look for non-managers with access to manager-specific applications." (p. 72)

Preventive and detective policy enforcement
"Detective controls are simply a periodic process to find undesirable states once they have already occurred. Good examples here are access reviews, reporting and analysis, and inventory variance assessment." (p. 73)

"If an organization has full preventative controls, do they still need detective policy evaluation? The simple answer is you need both. In the perfectly organized world of administration utopia, all changes happen in sequence and according to policy. In the real world of IT and business, things trend heavily toward entropy and errors, and omissions should be expected." (p.73)

Certification and access reviews

"**Certification and access reviews are an important part of the Identity Governance process.** They enable managers or other responsible delegates to review and verify user access privileges in a consistent and highly auditable way." (p. 73)

"Born out of corporate and financial audit requirements such as PCI and SOX, the process strives to force the business and IT security to come together to ensure least privilege and appropriate user access." (p. 74)

"**The Identity Governance server collects fine-grained access or entitlement data from all of its connected systems and formats the information intro structured reports that can be sent to the appropriate reviewers for verification.**" (p. 74)

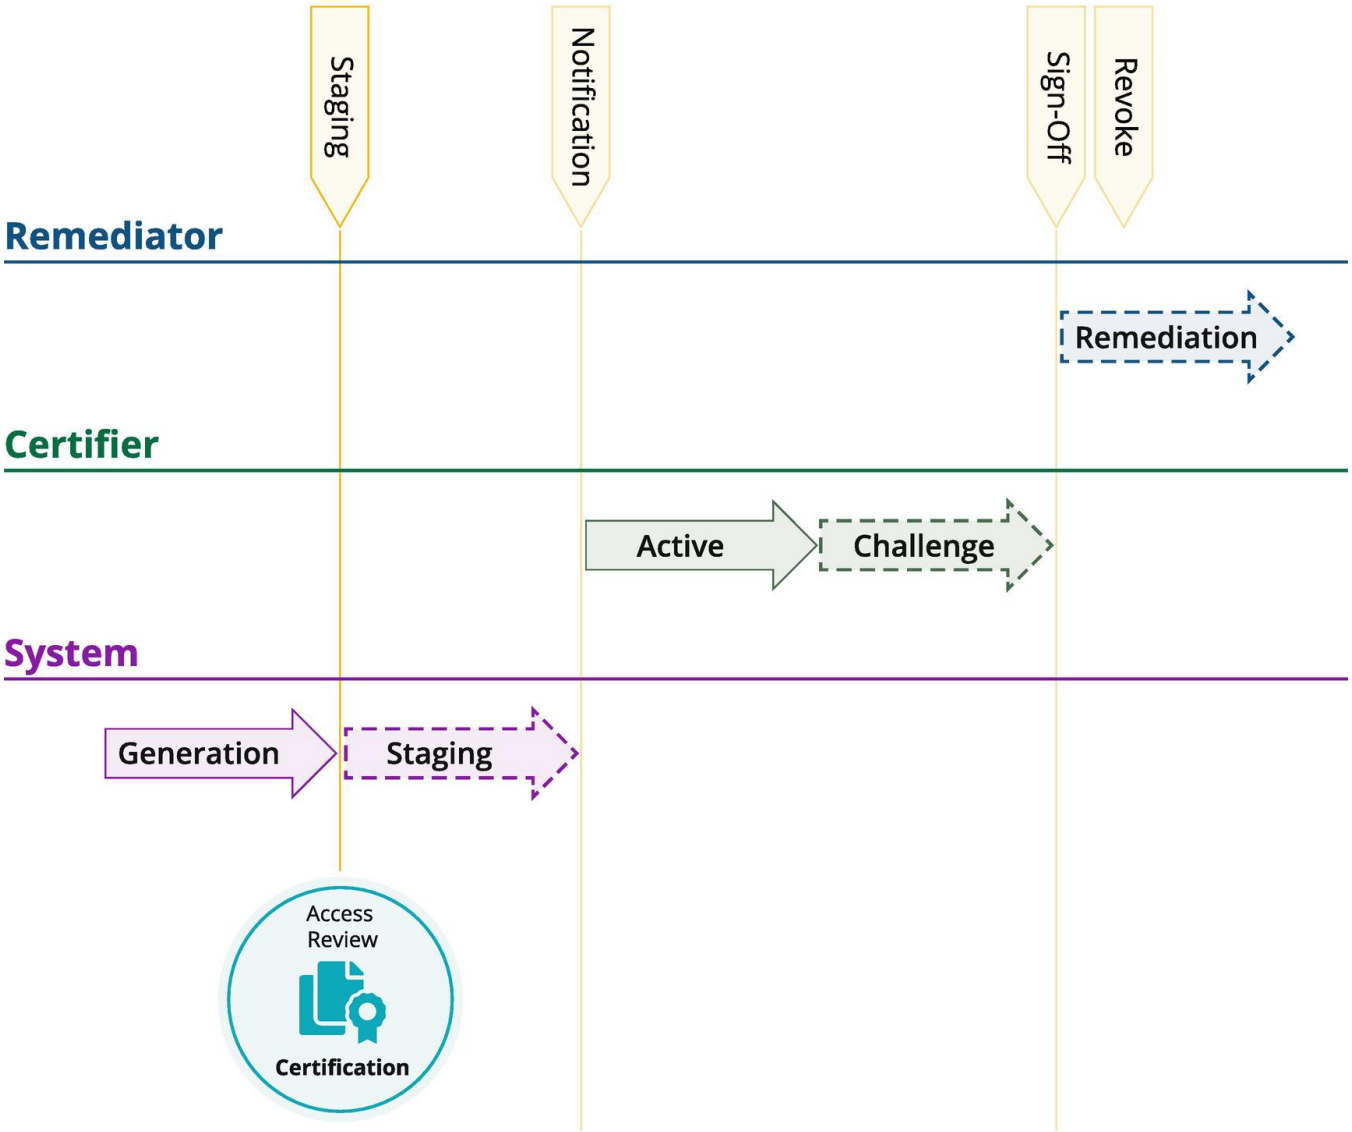

"The first step in the certification lifecycle is the **generation phase**. This involves specifying the dataset to be included in the review cycle and defining its schedule." (p. 75)

"**Staging** allows the system to generate candidate access reviews and stage them such that they can be checked before becoming visible to the certifiers." (p. 76)

The **notification phase** is "simply letting everyone involved know that they have work to do." (p. 76)

"During the **active phase** of a certification, the lines of business are actively reviewing and verifying access." (p. 76)

"**Some Identity Governance systems will also implement a challenge phase in the certification process. Here, identities are notified before a revocation affecting their**

entitlements is executed; this allows them the opportunity to dispute the decision and offer an argument for why they should retain the access." (p. 76)

"The **sign-off phase** is when all of the required decisions have been made for a given access review and the certifier is asked to formally close the review process." (p. 76)

"In the **revocation phase**, entitlements are changed in the source applications." (p.76)

Certification pitfalls

"The most commonly talked about pitfall in the certification process is the dreaded 'rubber stamp syndrome'. This is when an approver bulk-approves all access rights by 'selecting all' and clicking 'approve'". (p. 77)

"Certification fatigue can also be a factor in highly regulated environments." (p. 77)

"If possible, move to a role-based and delta-change based certification model in an effort to reduce the number of things that need to be reviewed." (p. 77)

This is something that could be included in our certifications, recommendations based on role, or peer-group analysis.

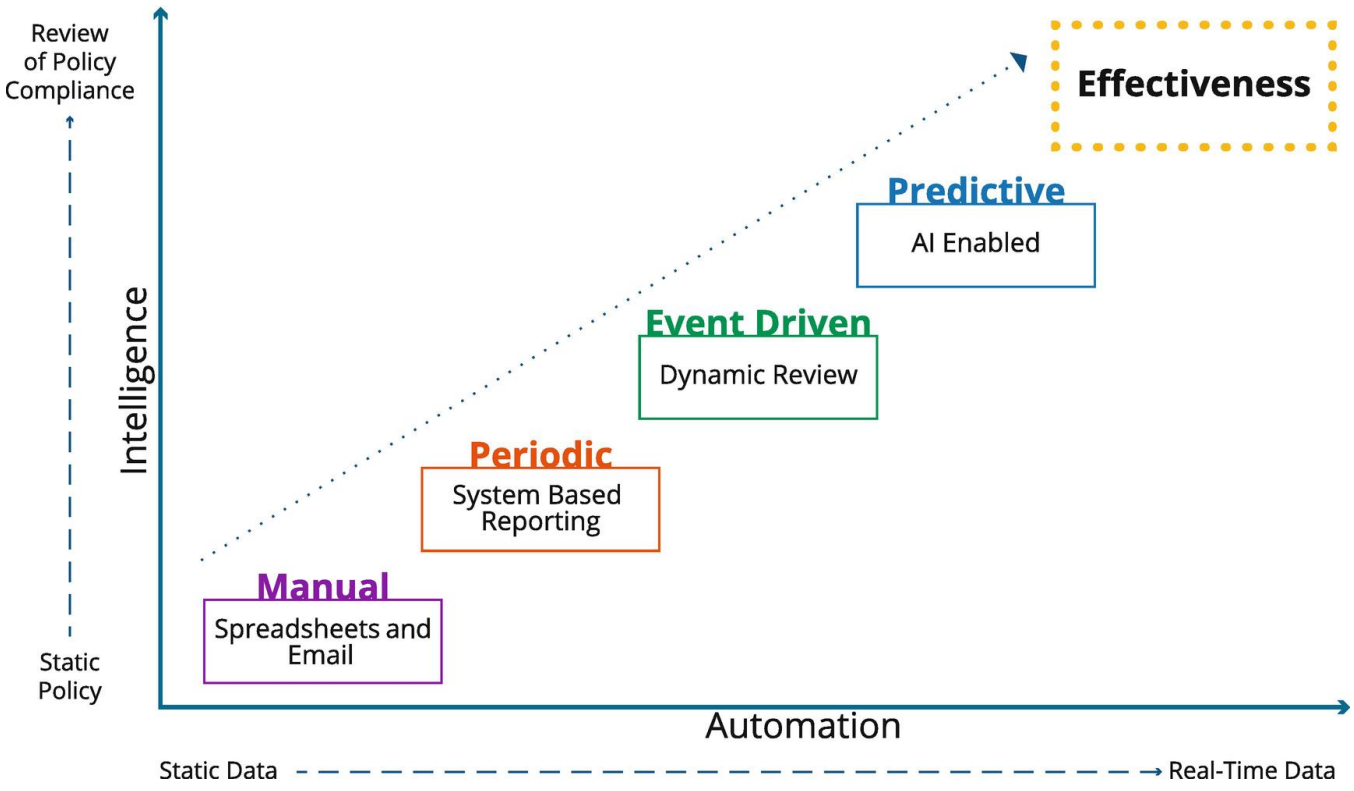

"Many organizations today still approach certification as a **periodic** control. They run bulk certifications on a regular rolling cadence, and everyone does the 'quarterly access review dance'. A periodic approach, delivered through a decent user interface that includes

current data and informed entitlement context, is more than enough compliance for many organizations." (p. 78)

"Many organizations are moving toward an **event driven** approach. For example, you might elect to review all administration rights for a group of admins on a quarterly basis, but then based on unusual admin activity, you might automatically re-review a specific admin." (p. 79)

"**The future of certification, however, lies in a predictive approach (ex. machine learning).**" (p. 79)

"**Here we see the governance platform making value and data-based decisions (based on behavioral baseline and peer-group analysis) to create a more dynamic and real-time approach to the process. Imagine your manager being asked to confirm your account groups for a mission-critical application, because you just logged in from an odd location - that's a responsive and highly predictive approach to a detective governance control and is rapidly becoming one of the key future drivers for leading technology providers in this space.**" (p. 79)

Enterprise role management
"The topic of enterprise role management could be a book of its own." (p.79)

"A well-designed RBAC system also simplifies and streamlines the administration of access, by grouping sets of access in a logical and intuitive way. Based on things like department, job function, title, persona or region, roles are assigned to users, and their access rights are automatically aligned with those roles out in the infrastructure. This provides a secure and efficient way of managing access and helps keep things simple for administrators, certifiers, and the users requesting access." (p. 80)

"Roles allow an organization to meaningfully move toward a 'manage by exception' paradigm by defining known groups of access aligned with business activities and functions and highlighting where the current state differs from the model view." (p. 80)

This is something we should communicate often.

"Quite simply, roles make certification and access reviews simpler, faster, and more business-friendly. During the business certification process, roles allow the user to focus on the assignment of groups of entitlements rather than getting lost in individual entitlement configuration." (p. 80)

"This vastly simplifies the administration and oversight burden and allows specialists to focus on defining and validating governance policies instead." (p. 81)

"Roles provide an important 'model construct'. A role (and its supporting metadata and control processes) provides a concrete model construct around which the business and IT can come together to define, capture, and enforce the desired state, hence helping to ensure that the right people have the right access to the right data." (p. 81)

Business roles
"A business role represents job functions, titles, persona, or responsibilities. They are usually tied to the organizational structure and are assigned to users based on their functions in the business." (p. 83)

IT roles
"It roles encapsulate sets of system entitlements. They are tied to actual permissions (which may contain privileged access) within an application or target system. They represent the actual state of the user's access, such as ab account, an entitlement, or set of permissions required to execute a given function." (p. 83)

Required relationships
"Required relationships refer to the set of access that someone with a given role must have. Someone with an accounts payable business role, for example, will always need to have read and write access to the accounting system." (p. 83)

Optional or permitted role relationships
"Permitted relationships refer to the set of access that is discretionary. These are groups of permissions or entitlements that a user may be allowed to have but isn't required to have. When optional relationships are used to connect an IT role to a business role, the entitlements defined in that IT role are essentially 'prescreened' - we know that a user with this business role is allowed to have the permitted access." (p. 84)

"It creates what can be thought of as a 'model-based least privilege' and can really help business and IT teams work together to better understand the lifecycle of access." (p.84)

Engineering, discovery and analysis
"Defining and validating a role really is an engineering exercise and an ongoing process, one that takes skilled practitioners, smart tools, and reliable infrastructure. An effective governance platform will provide a range of tools to help with the discovery and analysis process but will always be dependent on a skilled staff and solid understanding of the applications, network, and infrastructure environment in question." (p. 84)

Tools required for role engineering, discovery and analysis:

Entitlement analysis and search
"A frequently employed starting point for a role engineering project is entitlement analysis and search (ad hoc queries)." (p. 84)

Automated role mining
"**Role mining analyzes data discovered and collected in the Identity Governance system. It uses pattern-matching algorithms and peer-group analysis techniques to look for collections of, similar access and outliers.**" (p. 84)

"**In a two-tier role system, the Identity Governance platform should support automated role mining to create both Business and IT roles. Business roles typically model how users are grouped by business function, including functional hierarchies, project teams, or geographic location.** IT roles typically model how application entitlements (or permissions) are logically grouped for streamlined access." (p. 85)

"Business role mining facilitates the creation of organizational groupings based on identity attributes, for example, department, cost center, or job title." (p. 85)

"After the mining process is completed, the new roles are added to the system for lifecycle management." (p. 85)

Peer-groups and identity graphs
"**Peer-group analysis is an enhanced derivative of the classic discovery methodologies that is focused on building peer group data graphs and leveraging a broader dataset that often includes actual usage data.**" (p. 85)

Manual role creation
"**It is important that the Identity Governance systems provide an easy-to-use graphical way for roles to be input into the system, by hand or in batch.**" (p. 85)

Role lifecycle management
"Role definitions need to be carefully maintained and revalidated on a known control cycle." (p. 85)

"**We often see cases of an Identity Governance system being responsible for 'full assignment lifecycle and governance' when the full 'role model definition' is delegated to an external unrelated system. If you don't own the model, you don't own the integrity. Therefore, make sure your model is governed by your IG system and not something else.**" (p.86)

This describes the indirect connections between ISAM (application provisioning guidelines), and Approver Matrix and eSAF (IGA).

Enterprise role tips and tricks

- Take a pragmatic approach.

  "A comprehensive role solution could take months, or even years, to complete."

- Know what you're trying to accomplish.
- Look for groupings or role types.
- Enforce least privilege.
- Expect exceptions.
- Make roles reusable.

"If one person in the whole organization is assigned a particular role, maybe that access shouldn't be managed by a role at all." (p. 87)

Additionally, the concept of reusability also applies to the ability to nest roles within roles in a reusable manner.

- Involve the business experts.
- Test and verify your roles.
- Develop processes for role maintenance.

"Regular certification of role composition and role membership should be part of your ongoing program strategy. This should include a plan for how to retire roles when they are no longer needed." (p. 88)

The future of roles

"A traditional role model is an expression of 'expected access', across a known group of users (peers or even personas). The outliers are the entitlements and the people that fall outside of that group. This traditional, somewhat static view of the world is therefore greatly enhanced by an understanding of usage information and a more dynamic runtime connected dataset." (p. 88)

"**We see the future of enterprise roles enhanced by the construction of 'identity graphs' that represent relationships across identity, access, and usage information calculated as vectors. Using peer-group analysis and a range of graph algorithms, we can model scoped populations and carry out a far more granular and dynamic outlier analysis process. This process then results in essential new input into the role modeling process and forms the basis for a more predictive overall approach to the governance process.**" (p. 88)

Integrating ITSM and IGA self-service

Three basic options for integration of ITSM and IGA.

- **Launch-in-context integration**

  A requestable item is placed in ITSM and when the user selects it the IGA system is launched with the selection context maintained, preferably with SSO enabling a smoother transition.

  "It's accurate to call this integration model 'loosely coupled', but it's far from reasonable to call it 'highly cohesive'". (p. 92)

- **Pass-through fulfillment**

  "A core set of high-value access provisioning services are made 'first-class' service items in the ITSM platform, and the governance service is literally plumbed in as an unquestioning and ever-faithful fulfillment execution engine. A pass-through model usually means all the approval action and all the controls lie in the ITSM product scope, and its usual to see the Identity Governance provisioning fulfillment happen without workflow or any form of interactive approval in that system." (p. 92)

  "This integration model works well for exposing a small number of very well-defined services - often password reset and a small handful of new access provisioning actions. The challenges come from a lack of scope (a small number of service items) and the uncomfortable fact that having pass-through provisioning sort of violates the very purpose of having a single control point in Identity Governance in the first place." (p. 93)

- **Dynamic catalog integration**

  "A dynamic exchange of catalog items, approval processes, and fulfillment progress happens in real time. The Identity Governance team builds requestable units and marks them for publication in the ITSM catalog." (p. 93)

  "Both sides are able to track progress of requests as they move from approvals through fulfillment." (p. 93)

  "Obviously, this is a fully featured integration paradigm. Unfortunately, this level of integration is only provided by a small number of collaborating vendors and usually does not come cheap." (p. 93)

Managing requestable units
"We define a requestable unit as a service item with a well-known set of access entitlements and fulfillment requirements. Sometimes these requestable units are large buckets of access encapsulated in an enterprise role definition." (p. 94)

"**Searching and peer group recommendations - finding the right thing to request can be a challenge. An enterprise-class governance platform will provide extensive filtering and searching capabilities to help.**" (p. 94)

"**Providing peer-group searching and automated recommendations is becoming a key Identity Governance requirement. By leveraging artificial intelligence and machine learning algorithms, a next generation governance platform can provide real-time awareness to the access request process. Allowing people asking for things to see what others already have, or are requesting, provides key users insight and enables enhanced controls. Understanding outlier requests or suggesting likely peer-group entitlement recommendations vastly simplifies the request experience and allows for more dynamic preventative controls to be defined and implemented.**" (p. 95)

Delegation, capabilities and scope
"**Requestable units should be scoped such that only defined groups of delegates (or individuals) can see and request them.**" (p. 95)

Overlaying data-driven controls and governance
"**It's critical that the access request catalog is supported by dynamic and data-driven approval processes.** This means approval workflows should be driven by metadata on the requestable units and the applications and embedded in the delegation model, and not hard-coded in the approval workflow code. **This sounds like a technical nuance, but experience shows that getting this right can be the single most impactful factor affecting the long-term cost and maintenance of the Identity Governance system as a whole.**" (p. 95)

Chapter 8: Meeting Regulatory Compliance Mandates

"If you do nothing more than what's necessary to pass a SOX or FISMA audit, you are not likely to address your logical access risks or security requirements. Effectively managing user access risk requires meaningful diligence above and beyond 'checkbox' compliance." (p. 99)

Chapter 9: Indicators of Compromise

There are two forms of IoC analysis. "One that parses the entitlements for an identity, their associated accounts on potentially every resource, and documents their entitlements for reconciliation. The second is based on user behavior. It is an active analysis of an identity and their interaction with resources to determine the risk of their behavior. Essentially, this asks the question, was the user acting appropriately for their role?" (p. 103)

Best practices

- "Maintain a full map of accounts back to users using an Identity Governance system." (p. 104)
- "Minimize the number of service accounts associated with an identity." (p. 104)
- "Use a directory bridge to authenticate with one centralized authoritative source, like AD." (p. 104)
- "Avoid sharing credentials (especially administrator or root) with multiple identities." (p. 104)
- "An identity's entitlements should map to real user behavior in the privileged access management system." (p. 105)
- "User behavior should be mapped back to entitlements (bidirectional) to determine if behavior is appropriate." (p. 105)

Chapter 10: Identity Attack Vectors

"The Dark Web is nothing more than a collection of criminally inclined web sites that use service models to transact the purchase of passwords, tools, and data to leverage stolen information." (p. 111)

"Can you better attack the dead rather than the living? The simple and grace answer is 'yes'". (p. 112)

Individuals who are recently deceased:

- "Their bank accounts have not been closed or frozen nor their employer's ability for direct deposit."
- "Social media sites may allow active postings, and messages including ones used for business activities."
- "They probably still receive email at work or home."
- "Their cell phones and voice mail may still work."
- "They may not have loved ones immediately available to manage their estates."

"User behavioral analytics has a more difficult time interpreting malicious activity when valid credentials are being used and the account itself is already considered privileged." (p. 114)

Chapter 11: Identity Management Controls in the Cyber Kill Chain

"Forensic reports clearly show that identity management mistakes and weaknesses are the common faults in many breaches. These identity management mistakes and process weaknesses are things like poor account controls, weak passwords, orphan, dormant and rogue accounts, weak inventory of entitlements, and the over-assignment of user privileges." (p. 117)

Addressing IAM gaps

Weak inventory and cataloging
"An automated recertification can also be used to highlight escalation of privileges that tends to happen over an extended period." (p. 123)

Strong authentication
"With strong controls over sign-on, internal security teams can identify administrative access that happens at unusual times and from unusual locations." (p. 123)

"Strong authentication should always be used when logging into an Identity Governance system." (p.123)

Password controls
"Strong password policies make cracking and brute-force password attacks computationally time-consuming and costly for the attacker." (p. 124)

Lifecycle management
"Implementing strong JML state transition controls can help detect out-of-policy changes to the overall assigned entitlement model." (p. 124)

Pam governance

Request controls

File access governance

Chapter 12: Identity Management Program Planning

This section is a good outline of the goals, and implementation steps involved in a strategic plan. It is better read than summarized, and more useful as a preparatory step when a new program is under consideration.

Chapter 13: Privileged Access Management

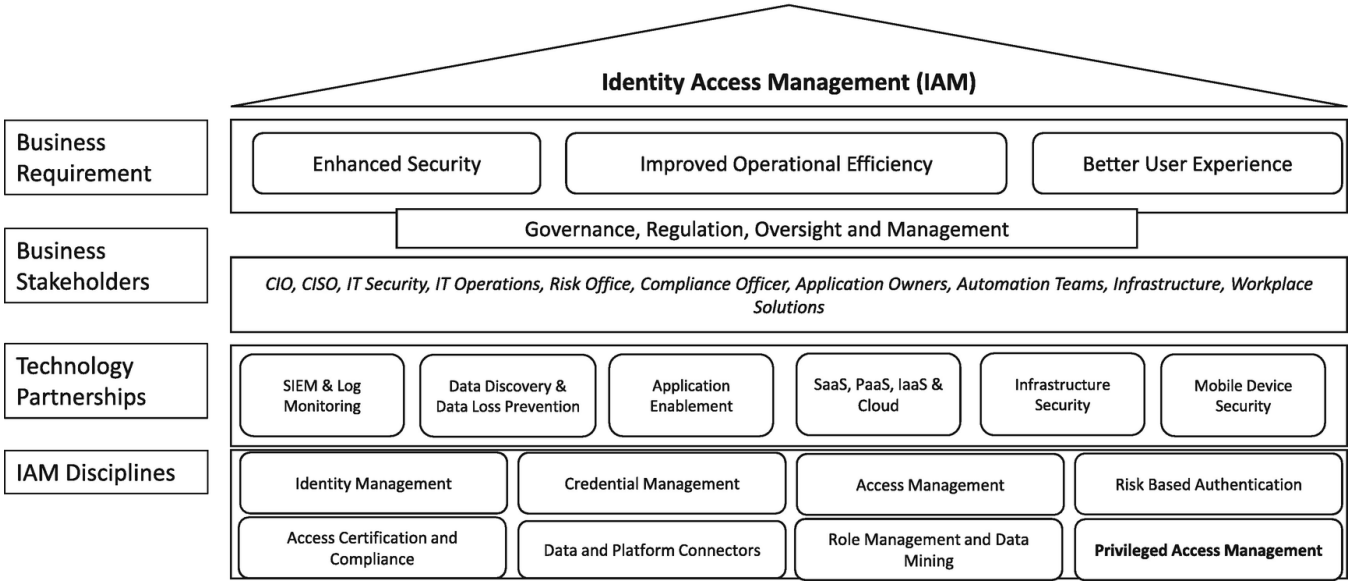

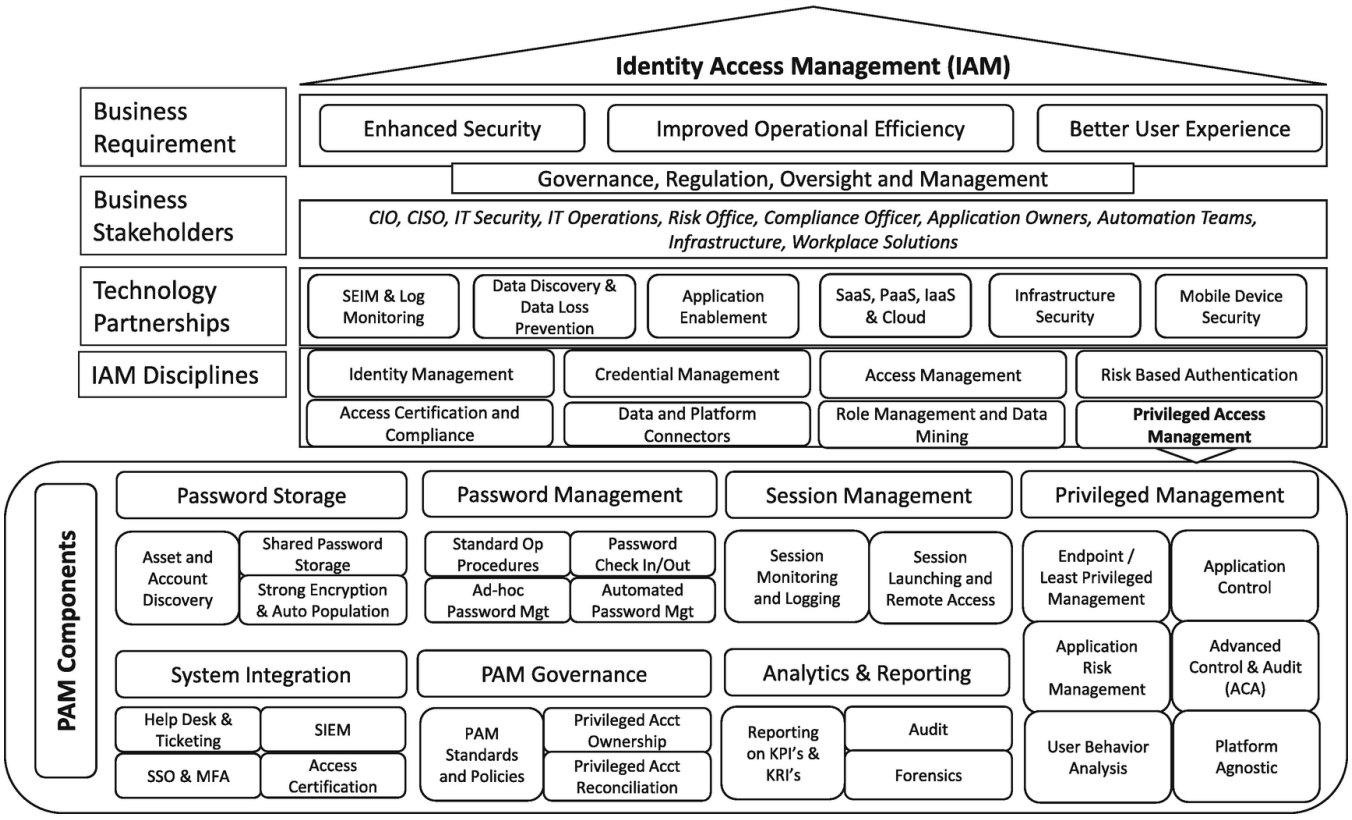

Chapter 14: Just-in-Time Access Management

"The concept of Just-In-Time (JIT) access management is a strategy that aligns real-time requests for usage of accounts directly with entitlements without the static assignment of an account or privilege to an Identity. Companies use this strategy to secure accounts from continuous real-time access by restricting them based on appropriate behavior, context, and other ephemeral properties. This decreases the risk of an always-on account that can be leveraged by a threat actor outside the acceptable use policies and procedures." (p. 151)

Key behavior attributes

- "Time and date windows for access and change control."
- "Commands or applications that may indicate a compromising event."
- "Detection of access to sensitive information."
- "Inappropriate modification of resources including installing software or modifying files."
- "Inappropriate attempts at lateral movement." (Observing login times)
- "The manipulation, creation, or deletion of user accounts or datasets."

"The risk of always-on accounts, unfortunately, is expanding. New highly entitled and privileged accounts are required for virtual, cloud, IoT, and DevOps environments in order to administer solutions." (p. 155)

Chapter 16: System for Cross-Domain Identity Management (SCIM)

"**SCIM uses a standardized REST API and data formatted in JSON or XML to allow interoperating solutions to exchange information in a standardized way.**" (p. 159)

"SCIM can also be used to share information about user attributes, attribute schema, and group and role membership." (p. 160)

Chapter 17: Remote Access

"Regardless of how you carve it up, the end user will still need an account to initiate connectivity. Whether you create the account in your domain or not may be a matter of your own security policy rather than the risk itself. The risk manifests itself in the access control, source of the connection, and connection and network connectivity required by the account." (p. 163)

In a remote connection, "the account may be trusted, but the asset (host) may not be trusted unless the device is issued by the organization." (p. 164)

"Therefore, the first step is to consider the account used for authentication. It should be under the control of the Identity Governance system and should have a real employee as the owner." (p. 164)

"Organizations should mitigate the risks from the asset and provide connectivity strictly by an account and with a least-privilege model." (p. 164)

"There is no way to know what resources have access to that account, especially if the asset connecting is not under your control. It could be a threat actor, trusted individual, untrusted computing device, and other combinations." (p. 165)

Session monitoring tools are useful in certification reviews.

Chapter 19: Biometric Risks Related to Identities

"Although many vendors seem to regard biometrics as a holy grail for authentication, large-scale breaches of biometric data and the inability to rest its source highlight a flaw in this approach. For these reasons, biometric data should only be applied for authorization - never for authentication alone." (p. 171)

"While there are systems that can validly claim to be impossible to hack TODAY, that claim only has a finite lifetime. Biometric identities are stuck with you for a literal lifetime. You cannot change your fingerprints or the blood vessels in your eyes. So, by the time we are older, there is a good chance that the system securing your biometric data will have been compromised." (p. 172)

"The OPM breach has shown that it is possible to steal biometric data." (p. 173)

Chapter 20: Blockchain and Identity Management

"Hopefully, you have not realized too late that they (blockchains) actually have a more limited place in business than the hype would lead us to believe." (p. 175)

"A blockchain is simply a distributed electronic ledger system that maintains multiple copies on multiple nodes. Blockchains are not a database replacement technology, they are simply a specialized computing technique that secures data entries based on a distributed system of cryptographic verification." (p. 175)

Cryptographic Hashing

Hashing

"A cryptographic hash is a mathematical algorithm for one-way encryption that is used to create tamper-proof on-line data. The crypto hash is at the center of everything relating to blockchain." (p. 176)

"It provides a one-way encryption process that is very costly to brute force and yet is very fast to validate." (p. 176)

"If I hash something and share with you both the encrypted data and its hash value, together, we can achieve something that is often referred to as 'verification without disclosure'. This ability to verify the integrity of something, without sharing it, is key to how blockchain-related systems work." (p.176)

Blocks and chains
"A block consists of some header information and a payload; the payload would be the data we planned to share, plus its hash value." (p. 177)

"The really important thing about blocks is that they live in a chain, with the structure and integrity of that chain provided by capturing the hash of the previous blocks. Should a malicious actor tamper with one of the blocks in the chain, it would invalidate that block's hash and break the chain." (p. 177)

Distributed consensus
"Distributed consensus is a very interesting concept and is one of the truly innovative elements of how Distributed Ledger Technology (DLT) works. First, you have to understand that there are multiple copies of the chain, each hosted by a separate node in the network. In the case of blockchain, it means literally thousands of nodes all maintaining an individual copy of the chain." (p. 178)

"Each node in the blockchain network is busy listening for transactions (payments), creating new blocks, and competing to be the first to publish that block to the rest of the network. And when a node wins that race and does publish a block first, it must also pass validation by the rest of the network participants. Unless a majority of the nodes confirm the cryptographic integrity of the chain, the newly published block is not actually accepted. Every node is, in effect, competing in a giant crypto processing race, with the express goal of being the first. This is the heart of how distributed consensus works. Literally, thousands of independent entities agree on the cryptographic integrity of the blockchain." (p. 180)

Proof of work
"Nodes don't just get to publish a block without concern; they first have to solve a cryptographic puzzle to go with it. This puzzle is called a 'proof of work'. In essence, a blockchain proof of work is an ingenious crypto guessing game, one that involves taking a block, adding some data to the end of it - a 'nonce', basically some meaningless text - and then generating a hash of the whole thing. This process is repeated again and again, until a specific numerical pattern is found in the randomness of the output." (p. 180)

"Realistically, this process of adding data, generating a hash, and looking for training zero data is a giant game of chance. It might take hundreds of millions of guesses before the right format output is randomly generated. This computational guessing game costs time and money. To be precise, it takes about 10 minutes to solve the puzzle on one of the largest dedicated server farms in the world. So of course, this computational exercise burns CPU cycles, which consumes electricity, which costs real money. And so as soon as a node does solve the puzzle, it quickly sends off its block for validation by the rest of the network. Remember, blockchain is a first to file and validate paradigm." (p. 181)

"But if a newly published block fails validation with the majority of the nodes in the network - the block gets ignored, and the publisher has wasted their investment in time and money

to generate it. The process of puzzle-solving and publishing of blocks is called crypto-mining." (p.181)

"Now you'd be right to ask yourself, 'Why do miners compete - why do they invest the time and money in a giant crypto lottery in the first place?' The answers are simple; the get paid in Bitcoin if they win. You might also ask 'Why does the blockchain system need miners - why set the puzzle and pay the price?' The answer is that for blockchain to work, it needs a community of participants. It needs multiple nodes hosting the chain, publishing blocks, and validating its integrity. For a system without the centralized authority to succeed, it needs to promote participation, in order to generate that much needed distributed consensus." (p. 182)

Limitations

"Once an entry is accepted by a majority of the nodes, it is considered permanent. Therefore, if you can attack the server, application, and ledger processes you can tamper with the blockchain with fraudulent insertions. This is how some of the recent cryptocurrency attacks have been occurring. The server and application have been the target, not the blockchain directly." (p. 183)

"Once an entry is made, it cannot be deleted, modified, or suppressed, just linked with a new entry that supersedes it. This makes blockchains most suited to new information and not to the storage of complex changing datasets." (p. 184)

Chapter 21: Conclusion

"We offer the following critical tenets of identity attack vectors and how identity access management should be embraced to make your implementation stand the test of time." (p. 190)

1. Think identity and not account

   "By centralizing data around an identity, enterprises have a single place to model roles, policies, privileges, and risk, a single place to build out compliance/audit policies, and a unified approach to provisioning, privilege management, and access control across the organization." (p.190)

2. Visibility is king! Silos are bad!

   "Everyone needs a central point of visibility." (p.190)

3. Full lifecycle Identity Governance is a requirement

   "It is critical to always manage the lifecycle of an identity and its access by tying it to the business policies and business owners that are responsible for it." (p. 191)

4. Deploy integrated privileged access management

   "It is essential that you take a comprehensive and integrated approach to the vaulting, auditing, and monitoring of privileged credentials and access. This means a lot more than choosing the right PAM vendor and the best technology. Specifically, integrating PAM and Identity Governance should be a 'must have', and providing audit, controls, and governance for the composition, assignment, and usage of the PAM infrastructure is of paramount importance." (p. 191)

5. Adopt a predictive approach

   "You should look to leverage this (machine learning) technology to help security staff, compliance teams, and the broader business user community make smarter and more informed access decisions. Specifically, look for knowledge, insights, and recommendations to be fully integrated with your Identity Governance process flows, delivering more real-time policies, better access certification decision, smarter controls, and more informed governance." (p. 191)

6. Implementing least privilege

   "With accounts and entitlements under constant attack, it becomes increasingly critical that we drive toward a least privilege approach to access." (p.192)

7. The user experience is everything!

"Having the right business user experience for these critical security processes is an essential part of achieving widespread participation and adoption." (p. 192)