Introduction to Access Control (v4)

By André Koot

To comment on this article, please visit our GitHub repository and submit an issue.

Table of Contents

Abstract

As the name implies, Identity and Access Management (IAM) is split into two functions: managing identity information and performing access control. Arguably, if there was no access control requirement there would be no need for identity management. It is therefore the focus for IAM professionals. At its core, access control is ensuring users are authenticated to access protected resources. This is accomplished by managing user entitlements and satisfying the requirements of relying applications so that users can only access the systems and information they are entitled to access. This article looks at the history of access management, the expected current functionality, and the trends to be expected.

Introduction

Access control, as a concept, has a long history. But in order to investigate the current challenges and solutions, let's start by evaluating a very old, traditional model of classified government documents.

Information in documents stored in files should not typically be accessible to everyone. The information may be classified, and only people with a required clearance level should be able to access classified files. In a physical form, this control is relatively simple: a folder with highly classified information is visually classified by a 'Top Secret' or 'For Your Eyes Only´stamp.[i]

But this simple example already addresses different fundamental concepts of security. First, there's the information itself. The information can be classified as Top Secret, but that must be defined by someone with the correct level of authority, like the owner of the information, the document, or the folder. Then, it must be clear what the impact of the classification level is; the classification level is needed to differentiate different levels of access to and usage of the information. The owner of the folder will probably have some guidance as to what levels of classification can be applicable and what type of user can get access.

Second, there is the clearance level of an actor, the user of the information. In this case, the secret service agent will have been identified and vetted to be trusted in such a way that access to different security levels of information is allowed.

Third, the classification level and the clearance level will have to be mapped in order to ensure that only the person with the correct security clearance level can access the classified information. The owner will classify a document and will accept that a specific security level can only be accessed by a pre-defined trust level of an agent.

And fourth: before giving the folder to the secret service agent, the person who is responsible for storing and retrieving the file in an archive (the file manager or access controller) must verify if the agent who requests the file is in fact the rightful user. The access controller will, therefore, try to identify the agent; the agent has to prove the right to access. This verification can be done by showing the secret service badge and a signed letter to prove that the agent has permission to access the folder. The file manager will, of course, also have to validate that the signature on the letter is correct.

Only after these responsibilities have been fulfilled will the folder be handed over to the secret service agent. The hand-over is then registered in a journal.

The access controller will always oversee the access, and that’s been made easy by checking the stamp on the folder. Theft of information—e.g., data leakage—is also quite physical in this example: the folder is removed. A folder with the 'Top secret' stamp should also not be found lying around unobserved.

In this scenario, access control is quite simple: you can literally observe access infractions. Access is granted by physically handing over the folder to a person with the corresponding clearance level, indicated by a personal badge, and may be enforced further by restricting access to a specific location.

We can see the following topics:

1. Classification of information: this is an aspect of risk management
2. Classification of users: this is an aspect of identity management
3. Authorization mapping: this belongs to authorization management
4. Authentication: this verification is part of both identity management and access management
5. Access granted: this is access control

Since the advent of the computer, there has been a need to control access to systems, documents, and other protected resources. In the early era of computers, processes analogous to the old spy movie era were used to model access control mechanisms. Concepts like ‘owner of a resource’ and ‘reader of a resource’ were used. Programmers developed access control mechanisms like Discretionary Access Control (DAC) (“you may never bypass the access controller,” a feature that can still be found in the Windows NTFS file system) and Mandatory Access Control (MAC) (“you can only access the data in a specific location” such as a dedicated workstation in a specific room).[ii,iii] The fast growth of information technology resulted in a growing need to develop and improve access control. The increase in the number of users, the number of systems, and the exponential growth of the information processed makes it evident that the paper world metaphor is not sustainable in the digital world.

It was soon realized that the concept of trust levels—e.g., managing the clearance level of an individual document reader—is hard to implement. Because so many actors are playing along and there is no longer a physical security control in place (you cannot see the red lint). Instead, there can even be multiple copies of a folder or file in multiple locations, and theft no longer means that the data is gone, but data will probably be copied without the consent of the owner. What was physically easy to implement is not easy to implement in the digital world. But the lessons learned in identification, authentication, authorization, access control, logging, and auditing, have been kept.

Access to information, data, services, and systems, as well as access to physical locations, is governed by security policies. These security policies must be formalized and need to be enforced by the owner of the resource. In doing so, the owner will try to manage the risk involved in access, such as the risk of abuse of information, data leakage, theft, fraud, and other security threats. In order to be in control, the owner needs to have the assurance of the level of security capable of being achieved by the security controls that have been put in place.

Apart from the concepts of access control, ownership in itself is a complex topic. Looking at the concept of data ownership, many criteria to establish ownership can be identified. Someone can be the owner of information because:

- They created the data.
- They funded the data processing facility.
- The data is about this person (e.g., a medical record of a patient)

There can be many more criteria to identify the owner, but this is part of data governance and out of scope for this article. In the case of medical files, the object, the patient, has several inherent rights to the data, making this person partly accountable for the access decision.

Terminology

- Identification – Uniquely establish a user of a system or application.
- Authentication – The ability to prove that a user or application is trustworthy and has the authority to access a protected resource by validating the credentials of an access requester (a user, a process, a system, or a thing).
- Multi-factor Authentication (MFA) – An approach whereby a user's identity is validated to the trust level required according to a security policy for a resource being accessed using more than one factor (something you know (e.g., password), something you have (e.g., smartphone), something you are (e.g., fingerprint).
- Authorization – Determining a user's rights to access functionality with a computer application and the level at which that access should be granted. In most cases, an

'authority' defines and grants access, but in some cases, access is granted because of inherent rights (like patient access to his/her own medical data).

- Accountability – The obligation of a person to accept the results of one's actions, be they positive or negative. This person is probably also a type of owner.
- Protected Resource - A system, process, service, information object, or physical location that is subject to access control as defined by the owner of the resource and by other stakeholders, such as a business process owner or risk manager.
- Access Control – Controlling who can have access to data, systems, services, resources, and locations. The 'Who' can be a user, a device or thing, or a service.
- Access Governance – The assurance that all access has been given based on the correct decision criteria and parameters.
- Access Policy – Definition of the rules to allow or disallow access to secured objects.
- Access Requester – The person, process, system, or thing that seeks to access a protected resource.
- Access Supplier – The component granting access to data, systems, and services after the access policy requirements (set in the Policy Administration Point) have been met by the Access Requester.
- Policy Engine - It is a security component that validates whether an actor is allowed to access a protected resource, following the requirements in an access policy. A policy engine can be seen as a component that exists of a PDP and a PAP combined.
- Policy Enforcement Point (PEP) – The authority that will only let an access requester connect to the access supplier if the Policy Decision Point allows it.
- Policy Decision Point (PDP) – The policy engine validates access requests and provides attributes against the access policy (as defined in the Policy Administration Point).
- Policy Administration Point (PAP) – The location where the different types of owners define the access policy.
- Policy Information Point – The authority that refers to the (external) trusted providers of attributes that will be used in the Access Decision. An example is the credly.com service that administers Open Badges of certifications, such as CIDPRO™ or the Certified Information Systems Security Professional (CISSP).

Acronyms

- ABAC – Attribute-Based Access Control
- ACL – Access Control List
- AIAC – Artificial Intelligence-Supported Access Control
- CBAC – Context-Based Access Control or Claims-Based Access
- CIAM – Consumer Identity and Access Management
- CRM – Customer Relationship Management
- DAC – Discretionary Access Control
- MAC – Mandatory Access Control
- PBAC – Policy-Based Access Control

- PAP – Policy Administration Point
- PDP – Policy Decision Point
- PEP – Policy Enforcement Point
- RBAC – Role-Based Access Control or (less frequently) Rule-Based Access Control
- ReBAC – Relation-Based Access Control
- SCIM – System for Cross-domain Identity Management
- SoD – Segregation of Duties

AAA: Authentication, Authorization, Accountability

Just as we showed in the classified document example above, in order to get access, a validated identity is key. The ideas behind this paradigm can be summarized by the concepts of AAA.

Authentication

Authentication is the process of proving that the user with a digital identity who is requesting access is the rightful owner of that identity. It can be as simple as using a password or as complex as providing a digital certificate. Both the Access Supplier and the Access Requester must be able to manage and consume the results of the authentication process.

Challenge - Response

The user might provide proof of this rightful usage by providing a secret that only the access requester and the access supplier know, like a secret code or a password. The underlying mechanism is called Challenge-Response. The Access Supplier challenges the Access Requester to prove his or her identity, and the subject will have to respond in the way the Access Supplier expects. The simplest way to do a challenge-response is by asking for a password or pin-code. But also, the CAPTCHA feature on many websites is a form of challenge-response: prove that you are a human being.[iv]

Knowledge – Possession - Being

But other than a CAPTCHA challenge, a known secret can be shared. It may not be sufficient to assure the rightful access because by sharing a password or by finding a password lying around (on a piece of paper, for instance), others may pretend to be the rightful owner. This weakness of the known-secret model means that the trust level of an access requester who uses just a password may not be sufficient for some applications.

After identification and even authentication, there is a degree of uncertainty in identifying the rightful owner, which should result in further evaluation of the level of access. A low level of confidence may be enough to give access to public information, but it will probably be insufficient to provide access to classified information.

Adding more proof of identity can be done by demanding more specific and unique identifiers. These more trusted authentication means cannot be easily copied or easily shared or stolen (it is not impossible, but the cost of copying a secure physical token can be too high to make it economically unsound to forfeit). In practice, this is done by introducing additional factors, such as tokens, certificates, and biometric proof. Requesting these additional proofs of identity can be requested either at the start of a session at the first authentication or during a session after a previous low-trust authentication has been found insufficient for getting access to a secured resource. In this case, the low-trust access can be enhanced by performing a 'step-up' authentication, requiring additional factors: the first step during login could be using a password, and then a second higher-level step could involve the use of a token or biometric proof.

Authorization

Authorization, often a synonym for the phrase access control, is the next step in getting access after the phase of authentication. It is the act of granting access to a specific resource, such as a computer application or a specific function within an application.

Authorization is closely related to the concept of authority. Someone, such as an owner, is accountable and, because of the ownership, is mandated to authorize others to access the protected resource. This accountability does not imply that the other person becomes the owner, but it does mean that several permissions, such as 'read' or 'delete,' can be executed. The owner stays accountable throughout the lifecycle of the data. Some of the tasks of the owner can be delegated to others in such a way that, for instance, a line manager may, within the boundaries set by the owner, grant read access to a resource to an employee.

Mainstream Access Control Methods

Currently, many organizations have security policies embedded in various applications, operating systems, and networking components. These controls are implemented in the form of Access Control Lists (ACLs), Roles, and DAC business rules. But these controls have to be designed and implemented in every relevant component. And these controls have to be designed in a consistent manner. If, for instance, a Segregation of Duties (SoD) restriction is defined for a specific process, every system, application, platform, app, and network component must support the SoD rule. If one of the many components is lacking SoD control, then the organization is not in control.

This decentralized implementation of security policies makes it challenging to implement centrally managed organization-wide controls. It is likely that not all controls are similar and that the security policy and conformity must be verified for every system or platform access request.

Modern Access Control

In modern implementations of access control, a policy engine is used to evaluate access policies centrally, and policy enforcement should encompass the ‘risk level’ evaluation. The business process owner, or data owner, tasked with managing access risk, will define the policies for which they are accountable. In some cases, there are multiple ‘business owners,’ and each is responsible for their part of the corporate security policy. This assignment of business owners can result in continuously changing access control policies.

There is much development in this area, with applications no longer maintaining the ACLs of users. Instead, they rely on identity management authorization systems that will, based on one or more access policies, make the decision regarding a user’s access request. Different stakeholders in a company are responsible for different policies. All applicable policies must be evaluated before access is granted. This method of fine-grained access control is a type of MAC.

Accountability

Accountability is a key responsibility in access governance. Making sure that every access decision is accounted for by an authorized person implies that ownership must be addressed. The owner must be informed about all activities under their control in order to be successfully accountable for the data under their stewardship.

Registering all activities in access control is an essential quality requirement. This record can vary in complexity from logging every authorization request (like granting or revoking authorizations or roles to and from people) to logging changes of authorizations within roles. The existence of this register is essential to be truly in control of access. The same is true for the identification and authentication process. There must be assurance from the part of the login mechanism, the operating systems, and the IAM solutions applied to make sure that every access request is validated.

Specific Access Control Considerations

Access control is not only a business decision. Other considerations inform how this activity must take place, including how users will engage with the control mechanisms as well as legal implications for what is (and isn’t) required.

The Human Factor

The user who needs to cope with the security controls can themselves be a roadblock on the path toward effective ’control.’ User experience (UX) is a critical success factor in every information security project. If the security controls are too strict, users may be deterred, or they may try to circumvent the control. This avoidance on the part of the user is often seen in consumer access: if a customer portal is not built with a focus on the user, then consumers tend to go elsewhere. That is a missed opportunity, resulting in low conversion rates. Consumer Identity and Access Management (CIAM) solutions are developed to prevent this behavior.

The lessons learned in CIAM are also being implemented in workforce IAM: UX is starting to make an impact. For instance, if a user accesses a company intranet portal from their home location regularly in a prescribed way, like using a VPN, the access control system could validate this behavior as a factor in the authentication process. It could decide not to require the repeated use of multi-factor authentication since it is a trusted user making use of a known, trusted connection; it's a well-known context resulting in better control of access.

Legal Implications

Access control has historically been looked at as a way to support business processes and is part of a larger information security and risk mitigation policy. The question of legal implications directly tied to access control practices varies from business to business, from sector to sector, and from jurisdiction to jurisdiction. There is no unambiguous answer as to the direct legal requirement for most access control practices as these policies are often woven into a larger program that is driven in part by any number of laws, regulations, or standards. Part of the role of an access control program or system is to ensure that it is flexible enough to support the larger risk management programs of the business or organization. In this way, questions about legal requirements and compliance implications can be addressed organically, allowing the organization the confidence it needs to operate and move forward.

In separate articles in the IDPro BoK, different aspects of laws and regulations will be illustrated in more detail.

Current state of Access Control

Mainstream Access Control Mechanisms

Several mechanisms support the implementation of access control. This section covers the more common ones: Access Control Lists (ACLs), Role-based Access Controls (RBACs), and Attribute-based Access Controls (ABACs).

Access Control Lists

Access control to a protected resource is based on the classification level of the resource. Every resource will be classified by the owner (or a delegated person) in order to define the security level of the resource. Based on the security level, security controls must be put in place to ensure the correct level of access. The access available, i.e., the permissions that can be granted, are also known as entitlements (fine-grained permissions to access resources). One of the earliest and best-known implementations of entitlements is by using ACLs In an ACL, the owner of the file defines what users can have what type of access: read, write, update, delete, whatever the owner accepts as usage. This concept is easy to understand and easy to manage for individual objects. And if the number of objects is limited, controlling access via ACL's can be enough. But when the number of users and the number of objects grows, ACL's can be a restricting factor.

Every owner of a file will need to define the ACL for the object. This distributed method of control implies that central control of access is non-existent. But, from an auditing perspective: it's relatively simple to find out who has access to a protected resource since that is registered in the ACL of the resource.

The concept of ACLs will be explained in a future article in the BoK.

Role-Based Access Control

Managing ACLs can be a tedious task. Managing access to resources on a user by user or entitlement by entitlement basis faces issues as populations grow. At some point, the issue of scale meant that a new access management approach was needed. RBAC is an approach of granting access to resources on a group level instead of on an individual level. In order to realize this, an intermediate component needs to be in place after that of the access controller. A role manager or a role owner has to be able to map the role of a user to an entitlement to a secured resource. This mapping looks easy enough, but in practice, this means that this person needs to work with different other responsible persons in an organization to make sure that the authorizations are not conflicting with the business processes and organizational structures of the organization. In the access governance article, this concept and the complexity connected with the governance model is further explained.

In the example of an internal company website, every company employee is made a member of a group called 'Company Employees.' The resource—in this case, the main page of the internal website—is secured in such a manner that access is granted only if a user is a member of this group. Another example is the line manager who can make a new employee member of the role account manager and behold, the access permissions connected to the role account manager, are available to the new employee. This non-individual oriented way of granting access makes managing access a lot easier.

A system owner can also create 'roles' within an information system to prevent the need for managing individual entitlements. The system owner of a Customer Relationship Management (CRM) system can define a role for 'customer manager' and group system authorizations (such as reading a customer record from a database or filling in a form) to that role.

In RBAC, we can identify a multilevel role model. On the one hand, we can identify the grouping of identities organizationally or hierarchically, defining organizational or business roles. On the other hand, there is a grouping of authorizations or permissions on an application function or platform level called system or application roles. Connecting organizational roles to application roles creates a very efficient way of granting and revoking authorizations. But it is also very easy to complicate authorization management by nesting groups: for instance, employees working on the service desk can be made members of the group ’ServiceDesk'. This group then could be made a member of the group Windows Administrators. By doing this, it will soon become hard to find out who has the authorizations of a Windows administrator. That would be not just the group of people who are members of the Windows Administrator role but also employees who are members of the role of ServiceDesk employee. This nesting can frustrate the insight by no small means; many IAM projects fail by the lack of un-nesting possibilities. Nesting also limits the auditability of RBAC environments; groups have to be un-nested in order to evaluate authorizations and potential conflicting authorizations.

Implementations, pros and cons, will be explained later in a future article about RBAC in the BoK.

Attribute-Based Access Control

ABAC builds on the RBAC model by introducing additional controls based on business logic. A major failing of the RBAC model is its static nature. Once an entitlement has been granted, it generally is always available to an end-user, until it is manually revoked. This longevity means that users wind up carrying access with them from role to role if proper cleanup actions are not taken. To address this, ABAC expands on the model, taking into consideration different characteristics of users and users’ attributes at the moment of determining if access should be granted. As a result, an access management system can make a decision based on the entitlements of a given user, as well as the time of day, the location of the user (e.g., on network or remote, geolocation based on IP address) the type of device (e.g., personal, organization owned, desktop or tablet), and other worker metadata. ABAC can be used both in real-time to control access at the time of the transaction, or passively controlling the assigned roles and entitlements based on user metadata. Both approaches require strong input and support from resource owners, Role managers, and people or organization managers to understand the needs of the user as well as additional support from analysts to help define the business logic.

For example: The Customer Relations Management process owner could define that everyone with the attribute ‘Business Role = Account manager’ can access the resource only if attribute ‘Allowed Time = defined office hours’. Multiple variations of this dynamic access control philosophy will be described later in a future IDPro BoK article.

The Future Direction of Access Control

Access Control by means of ACLs and RBACs is relatively static; the combination between a user and his or her authorizations are set and do not vary easily, and other authorizations require changes. But people move between jobs, change devices, change location, or get new tasks in a new context. Also, the risk level assigned to a protected resource can change because of a change in context or a change in applicable laws and regulations. Relevant changes may include:

- Extended organizations, internationalization, collaboration and federation, flexible workforce, meaning that in daily operations, people outside the scope of the traditional HR-operations may need to get access.
- Moving data processing to the cloud - leading to the development of new protocols, such as SCIM (System for Cross-domain Identity Management (the first time the acronym was used, it was called Simple Cloud Identity Management, I suppose this was deemed too simple or restricting ☺).[v]
- New privacy regulations, such as the GDPR.[vi]
- The usage of mobile apps, using modern protocols like OpenID Connect requires a flexible access control topology.
- Enforcing end-user consent and control - developments like User-Managed Access (UMA). [vii]
- Move to API-based access to micro-services - leading to new access management architectures based on protocols like OAuth2.

These restrictions and changes show that a more dynamic method for managing access is needed. The future direction of access control takes this into account, and various developments can be identified.

Dynamic Authentication

Access control is not a static event. When a user starts a session accessing services requiring a low-risk level, then identification with a username and password combination may be sufficient. When later on in the session, another trust level is required. For instance, when performing a transaction, additional identification, like a token, might be needed.

In order to adapt to these session dynamics, authentication in itself should also be a continuous process through, for example, the new concept of behavioral biometrics. Examples of changing needs for trust in the identity:

- User switches context (such as location). This switch could effectively place the user in another trust zone, and the session should be re-evaluated
- A user opens an email attachment, which by itself requires a higher trust level. This action should enforce additional authentication, such as Multi-Factor Authentication.

Adaptive authentication is a secure, dynamic, and flexible form of authentication. It enables validating multiple factors to determine the authenticity of a login attempt before granting access to a resource. The factors that are used for user validation can depend on the risk associated with granting a particular user access and may involve adjusting the authentication strength based on the actual context.

Policy-Based Access Control (PBAC)

A dynamic, flexible method is required for access control to become effective and efficient in extended organizations in collaboration environments with a flexible workforce. Policy-based Access Control (PBAC) is the paradigm to provide this flexibility. PBAC, also known as Claims-based Access Control or Content-based Access Control, takes some of the business logic introduced in the ABAC model and enhances it by layering additional context evaluation and dynamic step-up capabilities

The context of an access requester can change dynamically. The dynamic nature of policy management and enforcement could require step-up authentication within a session to cater for the higher trust level needed if the defined risk controls require it. A policy engine will be responsible for checking if the user attributes and context information at the time that access is requested, comply with the access policies defined by the owners of the security policies. Context information might include time of day, geographical location, or device type. The scalability of access is also enabled by making it possible to collect attributes from different trusted and pre-defined attribute providers. As an example: this person can access the Risk Management reports, but only if this person has the CRISC certificate. ISACA provides this certificate, so a lookup in the ISACA registry could answer the question regarding the CRISC certification (the mapping of the Access Requester to the ISACA member is out of scope for this discussion).[viii]

The central component in this architecture is Policy Decision Point, which evaluates access policies and returns a response to the access request. The Policy Enforcement Point then enforces the response either by code embedded in the application or, increasingly, via an API gateway. The Policy Enforcement Engine is a discretionary component in the access request flow.

As a further natural development, AIAC and ReBAC have to be mentioned.

Relation-Based Access Control

A new concept in access control is ReBAC, or Relation-Based Access Control. ReBAC addresses the possibility of making access control decisions using the relationship between the access requester and the other identities who can potentially be affected by the access control decision. These access decisions can be deduced from (amongst other services) social media network relationships of the access requester. An attribute such as ‘reputation’ can be evaluated and considered. ReBAC relies on the availability of large, distinct data sets (incorporating data from HR/Sourcing & Access/entitlement/behavior) and on AI to conduct the evaluations and recommendations for access decisions.

The direction for ReBAC is not yet entirely clear, and the development is not mature enough for mainstream implementation. We foresee the potential for implementation as part of predictive role mining technologies for dynamic ABAC implementations.[ix]

Artificial Intelligence Supported Access Control (AIAC)

We can expect much more in this area when we add the concept of artificial intelligence (AI). With a robust environment that classifies sensitive resources, it's now possible to take a sophisticated risk management approach to dynamic access control whereby the identity manager solution will alert on access requests that exceed normal risk levels. AI will also monitor access control requests alerting on out-of-normal activity. As such, it can be an addition to current RBAC and ABAC implementations. This concept is not yet mainstream, and we can hardly predict the direction, but AI and machine learning may add some value.

User Control and Consent

Privacy laws and regulations create a new awareness of access to personally identifiable information (PII). These laws and regulations have driven the concept of data ownership and consent by customers, employees, or patients. Data owners expect to be in control of their personal information, and in many cases, laws and regulations are mandating this. Several technological platforms have begun to spring up to fill this data ownership gap. Solutions like User-Managed Access, by Kantara Initiative, have made their way in the new access paradigms. Facilitated by the further development of protocols like OAuth, implementation of the concepts is made easier.[x]

Conclusion

Mainstream access control mechanisms like RBAC and ACL's have a long tail and will continue to have valid use cases in many organizations. However, as companies, governments, and organizations begin to require communications and collaborations outside of their traditional four walls, other ways of controlling access are required.

Mainstream access control methods are not able to deliver the growing need for flexible access control in a changing world. Modern access governance requires modern access control methods. There is a clear need for dynamic access control. Interestingly, the tools are becoming available, and implementation need not interfere with the current best practices: adaptive authentication, and PBAC can be added to an existing identity and access architecture. It takes some planning, based on a roadmap. And of course, it requires implementing elements of access governance.

Author Bio

André Koot is IAM and Security Consultant at SonicBee. His IAM experience comes from a financial accounting and auditing background. This background of anti-fraud detection and prevention business processes led to research in the area of authorization and access control principles.

Change Log

Date	Change
2020-06-17	V1 published
2021-04-19	Author affiliation change
2021-09-30	Updated definition for authentication
2022-12-15	V4 published: clarification to Policy Engine definition; minor editorial updates

[i] Wikipedia contributors, "Classified information," *Wikipedia, The Free Encyclopedia,* https://en.wikipedia.org/w/index.php?title=Classified_information&oldid=1120242140 (accessed November 24, 2022).

[ii] Davis, Shannon, "A Look at Discretionary Access Control," blog, TED Systems, 1 December 2020, https://www.tedsystems.com/look-at-discretionary-access-control/ (accessed November 23, 2022).

[iii] Rouse, Margaret, "mandatory access control (MAC)," TechTarget, December 2013, https://searchsecurity.techtarget.com/definition/mandatory-access-control-MAC (accessed November 23, 2022).

[iv] Wikipedia contributors, "CAPTCHA," *Wikipedia, The Free Encyclopedia,* https://en.wikipedia.org/w/index.php?title=CAPTCHA&oldid=1122595810 (accessed November 24, 2022).

[v] "SCIM: System for Cross-domain Identity Management," http://www.simplecloud.info/ (accessed November 23, 2022).

[vi] "EU General Data Protection Regulation (GDPR): Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation)," OJ 2016 L 119/1.

[vii] Kantara Initiative, "UMA Specifications," wiki page, last updated Jul 27, 2022, https://kantara.atlassian.net/wiki/spaces/uma/pages/29229182/UMA+Specifications (accessed 23 November 2022).

[viii] ISACA home page, https://www.isaca.org/ (accessed November 23, 2022).

[ix] "Data Mining and Predictive Analytics: Things We should Care About," Inside Big Data, 24 November 2018, https://insidebigdata.com/2018/11/24/data-mining-predictive-analytics-things-care/.

[x] Wikipedia contributors, "Classified information.”