White Paper

Access control to support security-minded sharing of information

Abstract

In an increasingly interconnected and automated world, we are dependent on the flow of information between people, organisations, and systems. An appropriate and proportionate control of access to information is essential for the safety and security of individuals, society, and nation states. Significant harm can arise from unauthorised access to, or modification of, information, and in some cases the inability to obtain access when required. The paper reviews current access control models and identifies shortcomings in respect of an increased need for information sharing. It considers the concept of entitlement to information, and the role of rights and obligations. Management of access to information in complex sharing situations is also discussed. The paper identifies several areas where further work is required to facilitate development of more sophisticated access control mechanisms that could better support current and future needs.

Prepared by Hugh Boyes on behalf of the 4D-SIG Security Working Group

Status: Final
Version: 1.0
1st October 2023

Contents

1. Introduction

The concept of controlling access is well understood in both the physical and digital domains. In the physical domain access control typically involves the use of barriers (e.g., walls, doors, fences, etc.), control tokens (e.g., keys, passes, etc.) and controllers (e.g., locks, card readers, guards, etc.) [Ferraiolo, et al.:2003, p.1]. In the digital domain access to systems and application is controlled using logical and/or physical measures. The logical measures typically involved authentication and authorisation of an account using one or more factors (e.g., passwords, 'PINs', one-time codes, tokens, etc.), where the account may be used by a person or enable machine-to-machine connectivity.

The common approach to access control involves a controlled asset and a controlling or managing entity. The controlled asset may be a physical location/space, a system, an application, a sensitive or valuable physical item, or information in a variety of forms (e.g., physical, digital, audible, or visual). The nature of the controlling entity will be context dependent. The term entity may apply to an individual, such as a person (human) or party (legal personality), or an actor (or agent) operating with a degree of autonomy within predefined parameters. The agent could be an application, system, or process running on a system.

Control, exercised by the controlling entity, involves establishing whether criteria for access to the controlled asset(s) have been satisfied. This model works effectively where there is a clear relationship between the controlling entity, the controlled asset(s) and the entity requiring access, e.g., a bank, bank account and bank customer. The situation becomes more complicated where the entity that creates the information asset(s) permits a third party to act as custodian or processor. In these circumstances control of access to the information is vested with the third party and no longer under the direct control of the originator.

Level of control can vary from a simple binary approach, i.e., access is permitted or denied, to more complex controls, particularly in the digital domain. The complexity arises where any information or functionality accessed in an application varies according to a user's role, or other specified criteria or attributes. An important difference between the physical and digital domains is that once digital information has been shared it can be difficult to control. If placed on the open Internet, then it is virtually impossible to retrospectively impose controls.

In situations involving role-based access control a user will be granted the ability to perform specific tasks or functions. The use of attribute-based controls is common where there are different classifications of information (e.g., government security classifications such as Unclassified, Secret, Top Secret). In these situations, in addition to any controls at an application level, access to the systems is usually limited to individuals holding the necessary security clearance. Adoption of this type of rigid control raises questions regarding the extent to which access control influences what we could or should do, versus what we need to do to satisfy business objectives. These classifications can be augmented by the addition of caveats, such as Official Sensitive, where the addition term is in essence a handling instruction rather than a state or classification indicator.

In seeking to control access it is necessary to understand the objectives or purpose of any controls. For example, through consideration of questions such as: what is being protected; against what or whom is protection required; who suffers consequential loss or damage; and what are the relationships between the information originator, the entity holding it and the potential user. In respect of highly regulated information such as personally identifiable information (PII), legislation requires consideration of the nature of the entity seeking access, the purpose of such access, whether the subject has consented to sharing or processing of the information, and what exemptions apply.

While there are some parallels between the control of access to information involving PII and intellectual property (IP), the latter typically relies on mechanisms such as licencing, non-disclosure and confidentiality agreements, to establish the legal framework for access to, and the sharing and use of sensitive information. Regarding PII, it is generally national authorities that address breaches of security relating to or misuse of PII. However, for IP, it is normally the IP owner's responsibility for seeking recompense for unauthorised access to, or use of, the affected IP.

Introduction of concepts such as the Industrial Internet of Things (IIoT), Industry 4.0 (I4.0), connected cyber-physical systems (CPS), and digital twins, has led to greater data sharing in a hyperconnected world. The relatively simple models of access control linked to assets under the control of a single entity are of increasingly limited utility. For example, where are the controls applied, by the grantor of access in respect or the requestor of information, and/or by the systems and processes using the requested information. This is a growing issue in situations where organisations' operational needs require federation of data, i.e., aggregation of data from multiple sources. Current practice seeks to control information by managing its circulation but has limited influence of how the information is used, or indeed misused, once circulated.

Development of more sophisticated access controls would facilitate use and processing of information across systems, under different ownership, and often in different legal jurisdictions. This white paper explores current and future requirements for more nuanced control in distributed multi-stakeholder environments when accessing federated data sources of varying sensitivity. It considers the concepts of entitlement, rights and obligation and discusses what requires authentication (i.e., a user or agent) and what is being authenticated to (i.e., the resource).

This paper is structured as follows: Section 2 reviews current access control models, Section 3 explores the concept of entitlement, while Section 4 addresses the needs of federated data sharing and complex distributed systems-of-systems. Section 5 identifies areas requiring further investigation and research and Section 6 presents the conclusions and recommendations.

2. Review of access control literature

In this section the term entity is used to represent the subject whose access is being controlled, where the entity may be a person, process, or device, that causes information to flow among resources or changes the system state [NIST:2006, p.3]. The term resource is used to represent an object that contains or receives information, and access to it potentially implies access to the data and/or information it contains [ibid.]. Resources include stored data and information, processes, programs, processors, peripheral devices, including sensors and actuators that are digitally controlled.

It is important to recognize that information is revealed by a combination of data items and their context. Without context, data can be meaningless, although through combination with further data a knowledgeable party may be able to add context and extract meaning from the data. This can complicate access control in circumstances where in isolation, or within a limited context, a data item is not sensitive, but when combined with other data in a data object (i.e., document, spreadsheet, etc.) it is sensitive. In considering the treatment of data objects, an author is not necessarily the owner, the author's employer, or the entity commissioning the production of the data object may be the owner.

2.1 Access control in context

A long-established view is that access control seeks to limit the actions or operations that a legitimate computer system user or programs execute on their behalf, with the objective of preventing activity that could result in a breach of security [Sandhu & Samarati:1994, p.40]. As such, access control is part of the suite of security services that are employed to determine the legitimacy of users and their actions or operations. For any entity (i.e., human, system, process) these services typically comprise:

- registration – the process of establishing the physical/logical identity of an entity seeking access to resources, including any enrolment and/or provisioning of credentials;
- identification – obtaining the credentials of the registered entity, e.g., a user providing their username;

- authentication – determining that an entity's claimed identity is authentic, e.g., through use of a password;
- authorization – determining what the entity is permitted to do;
- auditing – a deterrent measure involving collection and review of access control logs to identify unauthorised use, misuse of privileges or other suspicious activity.

Whilst authorization and authentication are fundamental to access control [Ferraiolo, et al.:2003, p.3], there is a distinction between authentication and access control [Sandhu & Samarati:1994, p.40]. Effective access control depends on the rigour of entity identification and the correctness of authorizations [ibid.]

Access control is typically based on a model comprising the following elements [Atlam et al.:2018, p.254]:

- *subjects: represents various entities that can be user, agents, or processes that make an access request to access system resources (objects).*
- *objects: describes system resources encompassing data or information that needed to be accessed by subjects/users.*
- *actions: represents various types of actions or activities that subjects can perform on a particular object such as read, write, execute, etc.*
- *privileges: these are the permissions that are granted to subjects to be able to carry out a particular action on a particular object.*
- *access policies: these are a group of rules or procedures that specify the criteria needed to determine the access decision whether granting or denying access for each access request.*

The effectiveness of an access control system may be expressed in terms of its safety. This is assessed in respect of system configuration (e.g., the access control mechanism and/or model) and the confidence that it prevents leakage of permissions to an unauthorized entity [NIST:2006, p.4]. A 'safe' access control system is one where no permission can be leaked to an unauthorized or unintended entity [ibid.]. Safety of access control is achieved through provable models or use of constraints [NIST:2006, p.33]. The latter can be a complex and difficult task depending on the model in use. For example, in Bell-LaPadula constraints are implicit [ibid.] whereas in RBAC proving constraints is difficult [Harrison et al.:1976].

From a business perspective, there is a need to promote optimal sharing and exchange of resources, while managing the risk of unauthorized disclosure or corruption of valuable information [Ferraiolo, et al.:2003, p.1]. If a sufficiently fine-grained access control mechanism is available this can enable selective sharing of information where in its absence, sharing may be considered an unacceptable risk [NIST:2006, p.3]. Depending on the configuration of the controls, an access control system can not only determine whether use of a resource is permitted, but also when and how it may be used [Ferraiolo, et al.:2003, p.2].

2.2 Types of access control

The DoD Trusted Computer System Evaluation Criteria (TCSEC or Orange Book) defined two important types of access control for digital systems: discretionary access control (DAC) [DOD:1985, p.14] and mandatory access control (MAC) [DOD:1985, p.21]. The term Non-Discretionary Access Control (NDAC) has also been used when referring to the group of MAC policies where rules are not established at the discretion of users, i.e., they are established through administrative action (NIST:2006, p.6).

2.2.1 Discretionary Access Control (DAC)

DAC restricts an entity's access to a resource where the level of access is controlled by the resource's owner [NIST:1994, p.26]. The resource owner may grant a set of permissions (e.g., read, write, execute) to individual entities or defined groups of entities, with access limited to only those entities and permissions explicitly enabled by the owner [NIST:2006, p.6].

There are several drawbacks to using DAC in an enterprise. For example, once an entity has been granted read access, information may be copied from one resource to another and there are no restrictions on information usage once it is received [ibid.]. From an enterprise security perspective, the protection and control of information usage is determined by individual resource owners rather than through the systematic application of system-wide policies or those of the enterprise [ibid.]. This is an issue where there are regulatory, legal, or national security considerations regarding the handling of information.

2.2.2 Mandatory Access Control (MAC)

Adoption of MAC involves access policy decisions being made by a central authority rather than individual resource owners, and these decisions regarding access right cannot be changed by the owner [NIST:2006, p.7]. This method is often prevalent where enterprise or system security policies specify that control decisions are not determined by the resource owner, for example, in respect of health records or classified national security information. In these situations, a labelling mechanism and interface constraints may be applied to enable the system to enforce protection decisions irrespective of the wishes or intentions of the resource owner [ibid.] Where a system involve information at different levels of classification or sensitivity, multilevel security models (e.g., Bell-La Padula Confidentiality and Biba Integrity models) may be employed to specify the access control policy [ibid.].

2.2.3 Role-based access control (RBAC)

The RBAC control model is based on assignment of entities to roles, within the system, application, or organisation, where access rights (i.e., permissions) are grouped by role names, and use of resources is restricted to entities authorized to assume the associated role [NIST:2006, p.7]. RBAC can be an effective and efficient way of establishing and enforcing enterprise-specific security policies [ibid.]. For example, separation of duties to prevent fraudulent transactions achieved by assigning mutually exclusive roles to individuals, i.e. no individual can both enter a received invoice and authorise its payment. Use of roles can enable efficient maintenance of security policies as organisations and processes change, enabling role permissions to be updated without the need to update individual resource permissions [ibid.]. In establishing role permissions, a least privilege approach is applied, restricting granted permissions to the minimum necessary to fulfil the resource's function. The use of role hierarchies enables roles to be defined with unique attributes, nesting roles where necessary to implicitly include the operations that are associated with another role [NIST:2006, p.8].

2.2.3.1 Temporal and workflow constraints

Role definitions may be subject to temporal and workflow-related constrains [NIST:2006, p.8]. Temporal constrains involve time-based restrictions concerning access to resources, such as, limiting access to specific business hours. In combination with workflow constraints, temporal constraints may be used to generate dynamic authorisations during workflow process [ibid.].

A business process, or workflow, typically comprises a series of activities within an organisation, each of the activities representing a defined task involving two or entities, for example, a user, a resource, and an action or decision. The process will involve dependencies between and required sequence of tasks. Within the workflow, the tasks may be performed by several entities (human or computational) in accordance with an established set of rules, i.e., policies and/or procedures. The organisation's policies will typically define specific roles, including the separation of duties, which need to be enforced using relevant access controls.

2.2.3.2 Chinese Wall policy

A variant of RBAC is the adoption of a Chinese Wall policy, typically found in financial, legal, and consulting organisations where there are risks of conflict-of-interest issues arising between different disciples or business teams. The policy was devised [Brewer & Nash:1989, p.206] to address commercial security issues and meet legal and regulatory requirements regarding the use of insider knowledge in financial institutions.

In a Chinese Wall policy [NIST:2006, p.10], company-sensitive information is categorized into mutually disjoint conflict-of-interest categories (COI), where:

- a company belongs to only one COI;
- each COI contains two or more companies;

- a COI includes similar companies, i.e., actual or perceived competitors;
- multiple COIs can exist, where each COI covers specific (e.g., sectoral) conflicts-of-interest; and
- the policy aims to prevent an entity accessing information for more than one company in any given COI.

Whilst this description of this policy deceptively simple, implementation and deployment are less straightforward [ibid.].

2.2.3.3 RBAC models

A role is a semantic construct that provides flexibility and granularity of assignment of permissions to roles and entities to roles [NIST:2006, p.17]. The access policy is formulated by assigning permissions to roles and making entities members of roles, whereby an entity acquires the permissions associated with the roles of which it is a member. There are essentially four types of RBAC model:

- core - comprises five administrative sets found in all RBAC models: entities (users), roles, permission, operations, and resources (objects), where permissions are composed of operations applied to objects [NIST:2006, p.17]
- hierarchical – involves the adoption of a role hierarchy, with an inheritance relationship between roles, thus addressing overlapping responsibilities and privileges [NIST:2006, p.16]
- static constrained – which adds separation of duties (SOD) properties as fixed constraint relations that are imposed on role assignment, e.g., to prevent an individual creating and authorising a payment instruction [ibid.], and
- dynamic constrained - like static constrained RBAC except the constraints apply on a session-by-session basis, e.g., may be able to create or approve a payment instruction but not perform both roles in respect of a single transaction [ibid.].

2.2.3.4 T-RBAC models

A proposed variant to RBAC is Task-Role-Based Access Control (T–RBAC model), where access rights (i.e., permissions) are assigned to specific tasks, and then the tasks are assigned to roles [Oh & Park:2003]. This differs from RBAC as permissions are assigned to roles but to the task being performed. This approach could address the SOD issues that are potentially inherent in the core and hierarchical RBAC models. However, it does introduce additional complexity and may still require the implementation of constraints.

2.2.3.5 RBAC limitations

The RBAC concept assumes that individual job functions can be neatly encapsulated in a set of permissions, which in practice is not a simple task [NIST:2006, p.19]. Contention arises between the need for safety (i.e., strong security) and the ease of setting up and administering the permission, this is in essence a conflict between the granularity of roles and the number required per entity. Within an organisation this is a trade-off between risk and the cost of administering and maintaining the access controls. Introduction of web-based applications, use of “X-as-a-service” offerings, and integration of enterprise applications complicates implementation of RBAC.

A significant issue is the means of achieving separation of duty controls for individual roles when using RBAC [NIST:2006, p.19]. This subtle problem arises where an entity has all privileges necessary to accomplish some critical function, where as a result the security controls are compromised regardless of role structure [ibid.]. For example, a system administrator typically has significantly elevated privileges to enable the individual to undertake tasks not available to most system users, and if such an account is compromised the unauthorised party would system level privileges. In this example the solution may require both logical and physical controls, the latter relating to the limiting administrator-level access to specific end user devices and/or locations.

2.2.4 Rule-Based Access Control (RuBAC)

The RuBAC concept is based on utilising pre-determined and configured rules to permit entities to access resources. However, as noted by NIST there is no commonly understood definition or formally defined standard for this approach compared to DAC, MAC, and RBAC [NIST:2006, p.20]. The term "Rule-based access" implies creation of some organization-defined rules, it therefore encompasses a diverse range of systems. For example, the rules in a firewall fall within this concept. RuBAC will assess all access requests, processing the using a set of rules that determine the rights/permissions of an entity to use a controlled resource as defined by a security policy. This approach can be used in conjunction with other access control models such as DAC or RBAC.

Rules can be developed to process a range of conditions, for example, security labels attached to resources, systems, and roles, as well as business derived labels. For example, an application could deny access to a new user account until the account has been activated using a link or code sent to the user's registered email account, thus fulfilling a business need to verify email addresses associated with a user account. As such rules are set by the organisation, they support use of MAC as they are not user changeable [ibid.].

While it offers flexibility in implementing and administering an organisation's security and business policies, unlike other access control mechanisms, it does not provide access assignments and constraints directly related between entities, actions/operations, and resources. In designing rules care must be taken to necessary constraints and permissions are implemented, hence RuBAC is often used in conjunction with other mechanisms [NIST:2006, p.21].

2.2.5 Attribute-Based Access Control (ABAC)

The ABAC concept is based on assessing an entity's requests to perform an operation on resources, which are granted or denied based on assigned attributes of the entity and the resource, and environment conditions, which are evaluated using a set of policies that are specified in terms of those attributes and conditions [NIST:2014, p.7]. The environment conditions are independent of the entity and resource, and based on detectable environmental characteristics, provide operational or situational context in which access requests occur [ibid.]. Examples of environment characteristics can include time, date, day of week of the request, location of the user (either physical or logical, i.e., an IP address).

Irrespective of the size or complexity of the system, implementation of ABAC requires the assignment of attributes to all controllable entities and resources, and the development of an appropriate policy encapsulating the required access rules [NIST:2014, p.9]. ABAC evaluates relevant attributes related to the entity and resource, considering environmental conditions and the allowable operations for subject-object attribute combinations as determined in the access control rules specified in the security policy [NIST:2014, p.8]. The core capabilities to evaluate attributes and enforce rules or relationships between those attributes are present in all ABAC solutions [ibid.]. Access privileges are indirectly specified through the rules that bind entity and resource attributes, which are typically expressed as either:

- Boolean combinations of attributes and conditions; or
- a set of relations associating subject and object attributes [NIST:2014, p.10].

In both cases the intention is to specify authorised or allowable operations, where the granularity of control is determined by the richness of available attributes [ibid.]

2.2.6 Temporal Access Control

This type of control includes a time-based element that restricts when authorisation can be granted. Thus, a subject may have access to an object for a particular interval or time, which may be related to the object's temporal characteristics. Control can take several forms:

a) Temporal Authorization Model (TAM) supporting time-based access control requirements in a DAC model [Bertino et al.:1998], where each authorisation has an associated periodicity constraint. The temporal restriction applies to the authorisation and does not take account of any temporal characteristics of the data or objects [Atluri and Gal:2002].
b) Temporal and Derived Data Authorization Model (TDAM) which instead of applying a fixed period as the temporal constraint evaluates a temporal formula to evaluate time-related conditions to derive authorization rules [ibid.]. Fine-grained control may be achieved by including relevant

attributes in the formula's design, for example, working hours and the calendar period for which access may be granted.

c) Temporal Role Based Access Control Model (TRBAC) which extends the RBAC model by permitting periodicity/interval constraints to be applied to enabling and disabling the role [Bertino et al.:2001]. Triggers define dependencies between changes of RBAC authorization states (i.e., enable/disable).

d) Generalized Temporal Role Based Access Control Model (GTRBAC) overcomes a limitation in TRBAC which only applies temporal intervals to changes in RBAC state. This model extends TRBAC by permitting use of interval and duration constraints on user-role assignment, role-permission assignment, and role enabling events. [Joshi et al.:2005]. Other features include allowing dynamic changes in the authorization states, accommodating temporal hybrid hierarchy and addressing separation of duty constraints.

2.2.7 Risk-Based Access Control

The traditional relatively rigid access control models discussed above, which are designed to apply static security policies. These models give the same outcome in different circumstances, providing the fixed access criteria are met. In contrast Atlam et al. [2020, p.104] review the concept of dynamic models that can adapt to environmental changes and unpredicted situations by accommodating contextual and 'real-time' information when determining the access decision. A comparison of the two approaches is provided in Table 1.

Item	Traditional Access Control	Dynamic Access Control
Features	It uses predetermined and static policies to determine the access decision	It uses access policies and contextual features that are collected at the time of making the access request to determine the access decision.
Grant decision	The access is granted only if it matches one of the rules in the access policy	The access is granted based on the context and the policy. The decision can be overridden based on the context.
Deny Decision	The access is denied only if it does not match any rule in the access policy.	The access is denied based on the context and the policy. The change in the context can lead to changing the decision immediately.
Examples	ACL, DAC, MAC, and RBAC are the common examples of traditional access control.	Risk-based access control, trust-based access control, and combination of risk with trust are common examples of dynamic access control.

Table 1: Comparison between traditional and dynamic access control approaches [Atlam et al.:2020, Table 1]

Adopting a risk-based dynamic approach, authorization decisions are made by determining the security risk associated with access requests, and weighing such security risk against operational needs together with situational conditions [Khambhammettu et al.:2013 p.86]. This approach is beneficial where lives are at risk, for example in healthcare or military situations [Atlam et al.:2020, p.104]. The risk-based approach is more adaptable providing greater flexibility and potentially offering the ability to handle previously unidentified threats. However, it is a more complex approach requiring careful selection of appropriate and effective contextual attributes and the development of risk-weighting algorithms.

In their review of risk-based access control models Atlam et al. [2020, pp.15-17] identified nine commonly used risk factors: benefits of user, action sensitivity, resource sensitivity, outcomes of actions, context, trust, risk history, access policies, and role. The most cited were in decreasing order: risk history (of the user), context ('real-time' and environmental information) and resource sensitivity. In terms of contextual factors, location and time were the most popular [Atlam et al.:2017].

As noted by Atlam et al. [2020, p.20] successful implementation of risk-based access control required an appropriate risk estimation technique. Of the papers they reviewed 18 out of 44 did not discuss a risk estimation process, a further 8 offered a mathematically approach where the equations are viable dependent and not readily transferrable to different contexts [Atlam et al.:2020, p.20]. For this type of access control to be implemented, a robust and repeatable risk estimation technique is required. Without such a technique it would be difficult to assess the safety and or suitability of any proposed model.

2.2.8 Spatial and Spatial-temporal access control

The advent of location-based services and their use by mobile applications has created a demand for spatially aware access control mechanisms, which consider a user's position within a reference space and the user's position in relation to the location of an object for which permissions are being granted. This type of access control takes several forms:

a) GEO-RBAC model, which extends the RBAC model by enhancing it with spatial- and location-based information [Damiani et al. :2007]. It employs spatial entities to model objects, user positions, and geographically bounded roles. User roles are activated based on their position.
b) Spatial-temporal access control model in which Fu and Xu [2005] proposed a logical framework that supports a coordinated access control model enforcing both temporal and spatial constraints. The approach is based on a Shared Resource Access Language (SRAL) for the specification of access patterns by a mobile device.
c) Location and time-based RBAC (LoTRBAC) model [Chandran and Joshi:2005], which builds on GTRBAC (See 2.2.6) incorporating a fine-grained spatial model including detailed location hierarchy and the notion of relative locations. In this model, time is uniform over all the entities of the RBAC model (i.e., users, roles, permissions) whereas location context (and any associated constraints) may be different for each entity.
d) Location-Aware Role-Based Access Control (LRBAC) model [Ray et al.:2006] which relates different components in the RBAC model with location and then employs location information to ascertain whether a subject should be granted access to the controlled object.
e) Spatial-temporal RBAC model [Ray and Toahchoodee:2007] which associates each component of RBAC with spatial-temporal information.

For cyber-physical systems and national infrastructure, incorporating spatial-temporal constraints into the access control systems offers a mechanism to dynamically manages access to sensitive operational data.

2.2.9 Context-based Access Control

The individual access control models reviewed in this section have been designed to handle authorisation decisions based on one or more criteria (e.g., role, rules, risk, time, or location). However, there is an increasing need to handle authorisation decisions based on the context in which access is required. This has been referred to as context-based access control [Fernandez et al.:2007], who offer several interpretations of context, such as:

- "Context is the location and identities of nearby people and objects and changes to those objects."
- A "logical set of resources accessible … depending on several factors. Some of these factors may include client location, access device capabilities, management policies of the access locality, subscribed services, user preferences, and level of trust."

They highlight kinds of context, physical and logical [Corradi et al.:2004; Gamma et al.:1994], and organisational [Kirsch-Pinheiro et.al.:2005] and proposed a unified context and context-based access control model illustrated in Figure 1.

2.2.10 Federated access management

For many organisations the need for access control relates to managing the use of their internal resources. However, for some organisations there is a business need to permit federated access to resources. For example, academic institutions providing access to academic publications for staff and students. Such access is typically achieved using access and identity management tools such as OpenAthens[1] and Shibboleth[2]. In essence these tools rely upon organisations (academic institution) enrolling users and subscribing to a service (access to publications) offered by publishers. When an enrolled user seeks to

[1] www.openathens.net

[2] www.shibboleth.net

access the full text of material that is not open access, the relevant tool is used to verify the user's identity as part of a subscribing organisation and existence of a current subscription by that organisation to the published material. This effectively extends a static control model across organisational boundaries, without the need for users creating accounts in each of the publishers they access.

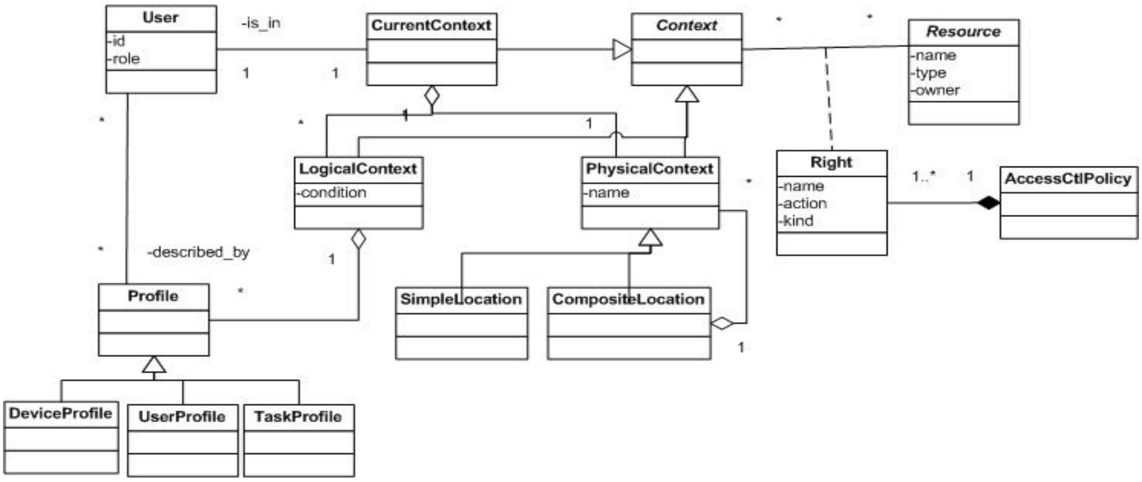

Figure 1 - Unified Context and Context-based Access Control Model [Fernandez et al.:2007, Fig.4]

A potential limitation of federated access management is the handling of volumetric control. For example, a publisher may not allow downloading of all the chapters from a book, or papers from a journal volume. This is unenforceable if two or more users from an organisation collude. Each may download content within the permitted volumes and then combine the content to reproduce the full work. Whilst this may be of limited concern to a publisher, if the data related to an organisation's intellectual property of the structure and location of critical infrastructure this may represent an unacceptable risk.

2.3 Limitations of access control systems

Current access control mechanisms are generally based on static control models and are designed to be managed by an entity (e.g., a single enterprise or organisation) to control access to its resources. There are exceptions such as the access federations aimed at satisfying publisher/subscriber business models. While use of static control models can satisfy access management in longer term relationships, they are not designed to handle ad hoc or dynamic requirements. A lack of dynamic access can be a significant hinderance where the information need is driven by spatial and temporal aspects. For example, in responding to a hazardous situation where emergency responders need timely access to sensitive information to which they are not normally entitled. In these circumstances access control based on a context aware model may be a prudent choice.

The use of federated access management tools offers some flexibility regarding use of resources by entities outside of the organisation controlling the resources. However, this is essentially an extension of the static models across organisational boundaries. For example, the requesting organisation determining which of its personnel may request access to the resources, and the resource publisher determining whether the requesting organisation has permission to access the specific requested content.

Current access control models are not designed to address the concept of entitlement, i.e., rights and obligations, or to address the enforcement of obligations. For example, if a user downloaded subscription content, whilst the user may agree to limitations on use there are generally not mechanisms in place to revoke access to downloaded content or prevent its onwards distribution. An exception to this is licenced content protected by a digital rights management application.

3. The concept of entitlement

3.1 Establishing the concept

The concept of entitlement can be defined as "*something that you have a right to do or have, or the right to do or have something*" [Cambridge Dictionary]. A practical example of this is holiday entitlement. In the UK there is a statutory entitlement for most workers to paid holiday, which is often referred to annual leave [Gov.uk], where a worker has the right to a minimum level of paid holiday. The concept of entitlement involves both rights and obligations. The former encompassing both negative and positive rights, and the latter addressing actions or behaviours to which an employee is legally or morally bound. Where an entity possesses a positive right (permitting an act or activity) there is an implication that another entity has a positive duty, for example, to take a specific action, and conversely a negative right (which obliges inaction) implies another entity has negative duties. In the case of annual leave, the employee has a right to a specified minimum level of paid holiday, and the employer has an obligation to permit the employee to take this paid time off work. In this case the rights and obligations may be codified by the terms of employment, limiting the duration of any single instance of annual leave, and requiring prior approval.

3.2 The nature of rights

Hohfeld [1913] considered that a 'right' was a legal interest that imposed a corresponding duty. He observed that the term was often misapplied to other legal interests (e.g., a power, privilege, or immunity) which were not strictly a right. Hohfeld identified eight fundamental legal (jural) concepts and two legal relationships which he considered represented all rights and duties. Figure 2 illustrates the concepts and relationships. The vertical arrows are jural correlatives, i.e., they represent a pair of legal positions that entail one another, while the diagonal dashed lines represent jural opposites.

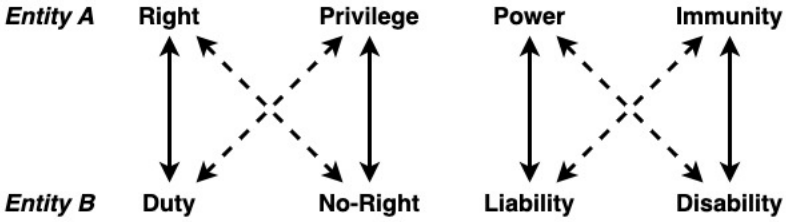

Figure 2 - Jural Correlatives and Opposites representing rights and duties.

Hohfeld proposed that the correlative 'obligation" limits the term 'right' to a specific meaning, where legal obligations accompany legal rights, and such rights can be legally enforced through claims adjudicated by the state. In Figure 2, the concept of privilege, relates to an entity's ability to act in a specific manner being liable for harm caused to other entities, who simultaneously are unable to seek intervention by the state.

Considering the four legal concepts on the right side of Figure 2, a power (or ability) is the opposite of a disability, where the entity with the power may exercise volitional control, for example, to vary the legal relations regarding a particular situation or problem. Hohfeld's concept of liability in this instance is best considered with respect to property or contractual transactions where a liability arises from susceptibility to an entity exercising its power. Just as power is the opposite of disability, immunity is the opposite of liability. An entity can be considered to possess immunity if it has independence from another entity's power, or control, over a (legal) relationship. A lack of power to alter entitlements in such a relationship is a disability.

A common conception of Hohfeld's rights and duties in the situation where Entity **A** owns or legally occupies a building or site, and only permits appropriately authorised entities to enter it, posting appropriate notices or signage at the entrances. Entity **B** has not been authorised to enter the building and therefore has a duty not to do so. This is a typical physical access control arrangement for government and military buildings and sites. Entity **A** has the power to authorise access to the building and exercises volitional control to grant access to some other entities but not to Entity **B**. The lack of authorisation implies that Entity **B** will have a liability if it enters the building without permission, e.g., the potential to be prosecuted for trespass.

Discussing the role of rights in respect of information, Risse [2023, p.77] refers to the need for both legal and moral rights. The inclusion of moral rights introduces additional scenarios that require consideration, for example, "how individuals behave or how institutions work in the absence of particular assignments of rights" [ibid.]. Another scenario he identified is the "right to know" [Risse:2023, p.78] in relation to medical ethics, where informed consent is contingent on the knowledge and if necessary, explanation of specific relevant facts. From a data protection perspective, the "right to know" personal data is addressed legislatively through mechanisms such as the ability to submit subject access requests and in some circumstances supplemented by the "right to be forgotten".

A further scenario relates to moral rights associated with the creation of knowledge. UK legislation [Copyright, Designs and Patents Act:1988] established the following four rights:

- to be identified as the author (the right of paternity),
- to object to derogatory treatment (the right of integrity),
- to object to false attribution, and
- to privacy in private films and photographs.

3.3 Epistemic rights

While Hohfeld's work was rooted in the physical world where 'rights' generally applied to the liberty of individuals, physical assets, or contractual services, the digital age introduces complexity related to the treatment of information. Risse introduces the concept of an "epistemic actor" defined as "a person or entity integrated into some communication network (system of information exchange) as seeker or revealer of information" [2023, p.106]. This definition is supplemented by several terms [2023, pp.105-6]:

- Knowers (or inquirer) – the entity that has acquired information, either by actively seeking it, or simply repeating information provided to it.
- Knowns – for individual knowns one can differential about what the entity reveals about itself versus what is otherwise known about the entity, e.g., through observation or inference.
- Revealers (or bearers of information) – entities that generate or disclose information, they may also curate it to preserve the information content in accordance with a set of societal norms and assigned roles.

Generalising Risse's work [2023, p.106], by shifting the focus from people to entities

- *individual epistemic subjects* – entities that gather and process information, abiding by standards rationality (i.e., seeking the best source of information) and moral, societal, and legal standards (i.e., who gets what kind of knowledge).
- *collective epistemic subject* – a collection of entities conforming to common standards of inquiry as individuals contribute to, or sustain, the information environment.
- *individual epistemic objects* – entities known by others delineated by rules concerning what information about itself may be shared.
- *collective epistemic object* – a group of entities that maintain and contribute to a pool of what is collectively known. The collective may determine the use and exploitation of the shared information.

Risse proposed that epistemic rights are "entitlements that justify performance or prohibition of actions, by the right-holder or other parties" [2023, p.109]. These rights concern which entities are entitled to what kind of information, where the entitlement may be described in terms of a privilege, claim, power, or immunity. Risse also proposed that epistemic rights are limited to the domain of inquiry, i.e., beyond learning of information X an entity may have no further entitlement regarding X, e.g., no entitlement to share X or exploit it [2023, p.111]. Epistemic rights are not the same as intellectual property rights, which relate to the economic use of information, as they address what an entity knows and how it is known.

To illustrate the application of epistemic rights, consider the following scenario. The employer of a system administrator (**S**) has commissioned a penetration test of their enterprise firewall by **T**. As an individual epistemic subject **S** is a knower and should be allowed to inquire about the test results. **S** has a *privilege-right* to know the result and no duty not to. **S** also has a *claim-right* against the penetration tester (**T**) to learn the result: **T** has a duty to inform **S** (and thus ought not to refrain from informing **S** or misinform **S**). **S** has a *power-right* to waive **S**'s *claim-right* and thus not know the test outcome, e.g., the result will be provided to a system engineer responsible for maintaining the firewall. Finally, an *immunity-right* protects **S** from **T** altering **S**'s entitlements regarding this information. There might be valid security reasons to regulate entitlements some other way. However, this scenario illustrates how the notion of an epistemic right operates for two individual epistemic subjects.

3.4 An example of rights and obligations

To further illustrate the role of rights and obligations, consider the operation of an entity's bank account and their entitlement. For an account that is in credit or within an agreed overdraft limit, the account holder is entitled (i.e., has the positive right) to make payments with the available funds. The bank has an obligation (i.e., a positive duty) to make such payments in accordance with the account holder's instructions in terms of payee, amount, and timing of any payment. The account holder has an obligation (in this case a positive duty) to ensure sufficient funds are available in the account to fulfil the requested payments. In operating the account, the bank is entitled (i.e., has a positive right) to refuse to refund payments made in error by the account holder or arising from the failure to protect their account credentials. The account holder has an obligation (in this case a negative right) to protect their account credentials (i.e., prevent disclosure) and to ensure the correctness of any payment instructions given to the bank.

The bank account example illustrates how in a relationship there can be complementary rights and obligations between the parties involved. Shue [1980] considered that all rights simultaneously involved both kinds of duty, thus respecting a right may involve avoidance (i.e., a negative duty) and protective or compliant actions (i.e., positive duties). He proposed that any distinction between positive and negative was essentially a matter of perspective or emphasis, i.e., a specific duty could be framed in positive or negative terms. Thus, in respect of a bank account, entitlements (i.e., rights and obligations) could be stated as:

- the account holder has the right, using their security credentials, to authorise the bank to make payments using their available funds and an obligation to prevent unauthorised payment instruction from being given to the bank; and
- the bank has an obligation to make payments in response to appropriately authorised instructions received in respect of an account, and the right to refuse to reimburse an account holder's losses in the event of disclosure of their security credentials, or where insufficient funds are available in the account.

To explore the concept of entitlement in more detail, two case studies are used, one concerning federated data and the other a complex systems environment.

3.5 Illustrating the entitlement concept – common information environment

In the context of federated information sharing, the concept of entitlement becomes more complicated as there are multiple entities involved. This is illustrated using a scenario, based on Figure 3 and described below.

A repository is established to share information about physical assets from numerous asset owning organisations. The data sharing agreements enable asset owners to specify the use cases (purposes), for which their information may be used, the type or nature of organisation that may access it, volumes of information that may be accessed, its granularity, and the permitted locations where it is processed, stored, or viewed. The shared information repository is managed on behalf of the asset owners by a custodian, who receives requests from registered users, authorises and manages release of information for processing, and reports on usage to the asset owners. A registered user can browse an index of information held in the shared repository, request access to and processing of the information and receive results from the information processor.

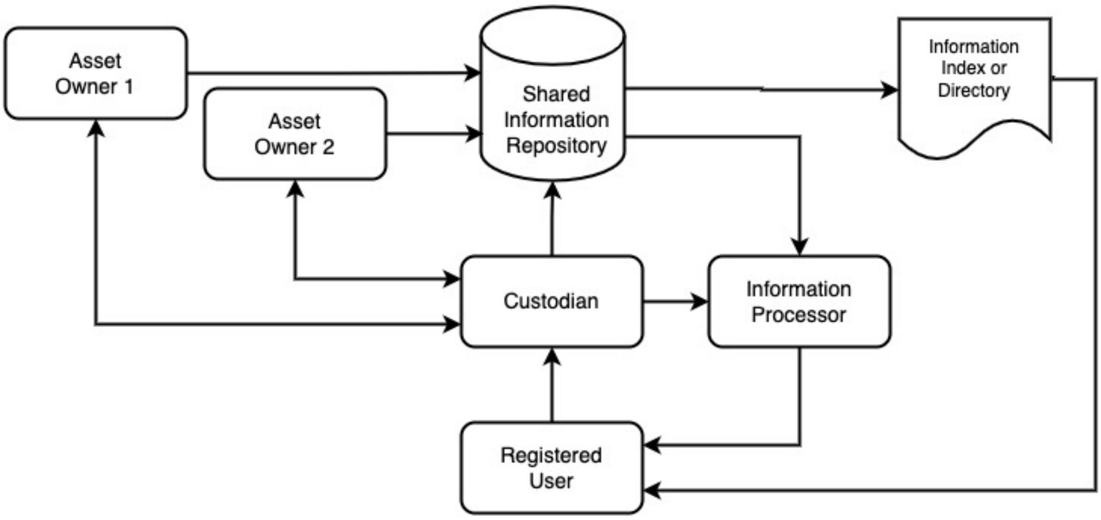

Figure 3 - Information sharing example

To establish a registered user's entitlement to access and process the requested information, the custodian has a duty, or obligation, to considers whether:

a) the request from the registered user falls within one of the purposes permitted by the relevant asset owner(s);
b) the registered user's organisation was of a type that was permitted to access the requested information;
c) the requested data volume(s) and granularity were within the limits specified by the asset owner(s); and
d) the location(s) in which the data was to be processed and/or stored were acceptable to the asset owner(s).

In this scenario access to the information is not simply a case of a registered user logging in to the system to use an application operated by the custodian. A registered user's entitlement is based on their request for information meeting a combination of specified criteria. Furthermore, the actual processing of the information must meet specified constraints, for example, those related to information sovereignty and/or retention. In a federated information sharing arrangement, consideration will need to be given to the management and protection of intellectual property rights, and the treatment of derived information. This may be particularly problematic where the processing includes AI/ML tools that may be able to infer information that asset owners did not intend to share.

3.6 Illustrating the entitlement concept – complex systems-of-systems

The Industrial Internet of Things (IIoT) and concept such as digital twins involve distributed data collection, processing, analysis, and visualisation, where individual systems, or sub-systems, may belong to different organisations offering functionality as a service, e.g., infrastructure (IaaS), software (SaaS, processing (PaaS), etc. Access to information may be subject to complicated contractual provisions, particularly where elements of the process are in different geographical regions and/or legal jurisdictions. The entitlement of individual actors/entities would be governed by operational, contractual, and legal rights and obligations.

A scenario is illustrated in Figure 4 where, for example, a transport infrastructure operator employs several service providers to provide enabling operational services (Application 1), some physical and others digital. As part of the digital service provision, an information service (Application 2) is used to inform third parties, e.g., end users and other infrastructure operators, of the status, availability, and current operating schedule for the transport service.

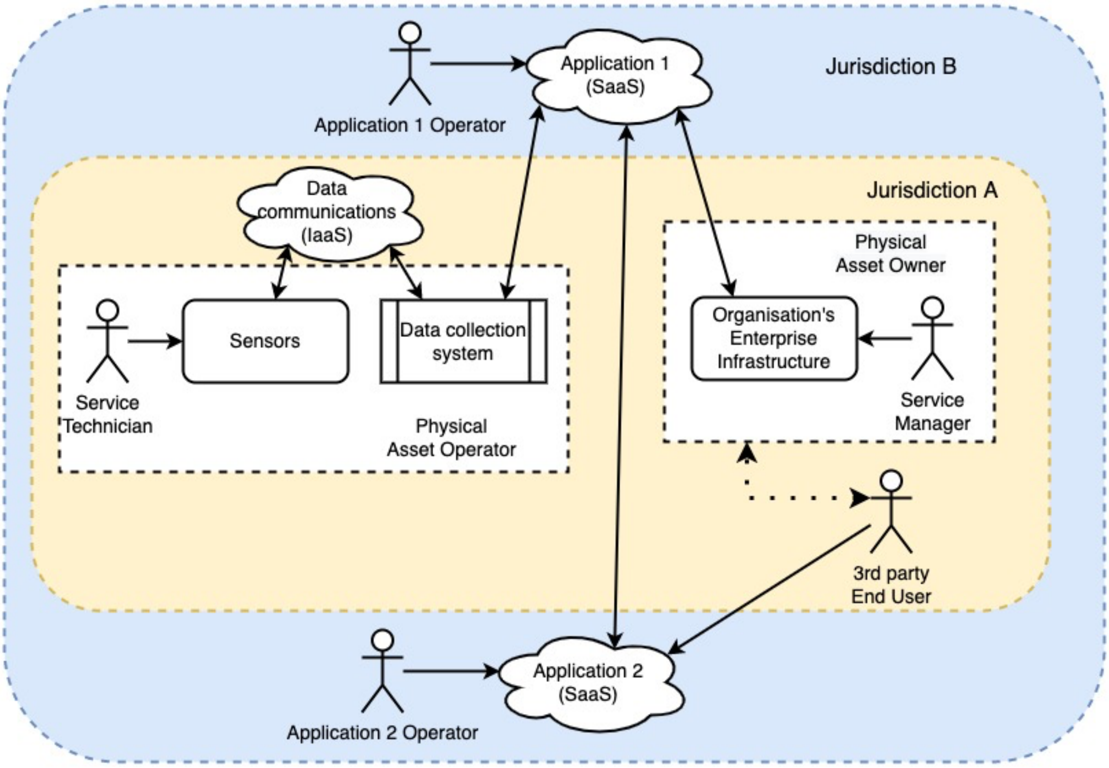

Figure 4 - Entitlement in systems-of-systems scenario

Table 2 illustrates the high-level entitlements of the various service participants in respect of the operation of the transport service.

Table 2 – Example entitlements for participants in Figure 3

Participant	**Entitled to**	**Not entitled to**
Physical Asset Owner	• Recruit and manage asset and service management personnel. • Schedule work of its personnel. • Manage contracts with service and application providers with supporting performance and financial data.	• Access personal data regarding Physical Asset Operator's personnel. • Access personal data regarding 3rd party end users.
Service Manager	• Create, read, update & archive asset data. • Monitor performance of assets. • Schedule use and maintenance of assets. • Monitor performance of application operators. • In an emergency have access to relevant personal information regarding 3rd party End Users.	• Access personal data regarding 3rd party end users. • Access personal data regarding Physical Asset Operator's personnel (e.g., service technicians).
Physical Asset Operator	• Recruit and manage service technicians => access to relevant personal data regarding technicians. • Schedule work of service technicians.	• Access personal data regarding Physical Asset Owner's personnel. • Access personal data regarding 3rd party end users.

Continued …

Table 3 – Example entitlements for participants in Figure 3 (continued)

Participant	Entitled to	Not entitled to
Service Technician	• Create, read, update & archive asset data relating to sensors. • Receive work instructions/tasking. • Report on completed tasks.	• Access personal data regarding Physical Asset Owner's personnel. • Access personal data regarding 3rd party end users.
Application 1 Operator	• Manage data exchange and integration between Asset Owner and Asset Operator. • Hold and curate data to be supplied to Application 2 in respect of service availability, operation and/or performance.	• Access 3rd party end user data held in Application 2, except where it is being passed from Application 2 to Application 1 with regards to handling of accidents, service emergencies, and user complaints. • Access personal data relating to Asset Owner and Asset Operator personnel.
Application 2 Operator	• Managed access to 3rd party end user data for the purposes of operating the service and providing user helpdesk. • Have read access to data held in Application 1 for the purposes of operating the Application 2 service. • Release specific end user information to the Service Manager from Application 2 via Application 1 regarding accidents, service emergencies, and user complaints.	• Create, update or archive data held in Application 1. • Use 3rd party end user data for the purposes other than operating Application 2. • Have access to information about the operation and ownership of the assets other than that provided to it via Application 1.
3rd party End User	• Create and maintain a user account for Application 2, e.g., containing user contact preferences, favourites, etc. • Access service timetables, planned outages and other data relevant to the use of and access to the service by 3rd party end users.	• Access Application 1. • Access personal data regarding service operators/managers. • Access sensitive information about asset and service configuration, maintenance, or performance.

4. Managing access in federated information sharing and complex systems of systems

In the previous section, the examples discussed may be treated as bilateral arrangements where the information creator makes data available based on a service contract, e.g., between a physical asset owner and an application service provider. The 3rd party end users are granted access via the end user agreement associated with the application, or website, that disclosing information about the transport service. A clearly defined architecture with access control defined by the transport service operator.

In contrast, a federated information sharing architecture could be an evolving ecosystem of information sources, where information creators permit access to their data with varying levels of discoverability, for a range of purposes. In a federated information sharing arrangement, each information creator retains autonomy and responsibility for their information, for example, its quality, frequency of update and continuing availability. However, as part of the federation arrangement the information creator will cede defined responsibilities to the entity managing or operating the sharing arrangements.

These sharing arrangements are likely to take one of the following forms,

- a searchable directory of information resources, which may be open access, or prospective users may need to request access from the information creator, or access information hosted on storage provided by the directory operator (custodian);
- a central repository, generally for a specific common purpose, into which the information creators place their information and from which authorised users retrieve it; and
- an information integration and exchange approach that enables users to search for and retrieve information that the creators have either explicitly or implicitly authorised their use.

The following sub-sections examine these forms and potential associated access control requirements.

4.1 Searchable directories

These are typified by the "open data" repositories created by public authorities, for example, the UK public bodies open data[3] and GLA's London Datastore[4]. There are typically few, if any, controls over who can access the information, the presumption being that it is openly shared. Its use and attribution may be governed by licences at the level of individual information sets. Where a information creator has specific licencing needs, for example with regards to licencing for commercial use, a link or email address to permit users to request the data may be provided.

Whilst the use of open searchable directories enables discovery of information, it is problematic if some information sets are sensitive. This is issue if knowledge of their existence is likely to attract unwanted attention. Where a directory will contain reference to information across the ODI spectrum (i.e., from open to closed) consideration should be given to limiting visibility of directory entries to the to the classes of information that a user is able to access, or exceptionally, is likely to be able to access.

4.2 Central repository

An example of central repository is the sharing of underground asset information for the purpose of managing safe digging during street works. Ownership of information about individual underground assets remains with asset owners, who collectively agree to share a specified information set. Individual asset owners may choose to exclude some asset information, for example, where the assets are at sufficient depth, they are unlikely to be damaged by normal street works, or where physical control of access limits the risk of unintentional damage. This type of repository may be used for a single common purpose, or multiple defined purposes, and be capable of supporting display of geospatial data from multiple public authorities.

Access to the information is typically via a dedicated application where results can be rendered in an appropriate format based on the relevant use case. For example, the display of underground asset information may support street works, revealing assets in a tightly defined area, as well as an asset planning mode where managers can plan maintenance and upgrades over a larger area. Depending on the permitted use cases there may be varying levels of access control linked to the risk of unauthorised use or disclosure of the information. For example, individuals capable of:

a) managing the overall system and its technical configuration;
b) setting up and removing user organisations;
c) administering users access within their own organisation;
d) undertaking asset planning activities; and
e) accessing data as part of on-site street works delivery.

[3] https://www.data.gov.uk/
[4] https://data.london.gov.uk/

4.3 Information integration and exchange approach

This approach may be employed to support cross-sectoral data sharing where information is to be combined to support specific applications or decision-making. An important difference between this approach and the use of a centralised repository is the nature of the retrieval. The central repository may employ a single standard information model for its content. This makes it easier to retrieve and combine information from multiple sources. In contrast the integration and exchange approach functions as a mechanism to retrieve information in varying formats, levels of granularity and aggregation. This can enable information creators to serve appropriate information according to the purpose for which it will be used.

From an access control perspective this is a more complex access control situation as information creators may wish to manage availability of information depending on the intended use and factors such as the user, information coverage, and the levels of granularity and aggregation. For example, where the information may be of value to a competitor, an asset owner may wish to limit access to infrastructure capacity and asset condition information. This is an application where more sophisticated granular access control may be beneficial or for some asset owners a commercial necessity.

4.3 Open Energy – Access Control and Capability Grant Language

This language [Open Energy:2022a, 9] is proposed to support data access by "*determining what types of conditions may be specified for {data consumer} to meet in order to gain access to datasets in different sensitivity classes*" [Open Energy:2022b, 3.4]. The language specifies how access rules are articulated in terms of a syntax and the conditions for granting access.

Using the language, access rules, which are unary or binary, are composed of:

- Zero or more conditions for grant access; plus, if access is granted,
- One or more capability grants to the data consumer; plus
- Zero or more obligations falling on the data consumer. [Open Energy:2022a, 9]

The rule syntax is a significant constraint, it can support multiple conditions, but all must be met to grant access, and the use of sub-clauses or Boolean operators is not permitted. The standard capabilities that can be specified with regards to information are specified in four groups: internal uses by the information consumer (use, adapt, combine) and where applicable permission to redistribute (share). There is some finer control within these groups [Open Energy:2022a, 9.3.1] but it is not clear how these would be enforced. While the developers of the language refer to obligations these are primarily in respect of what an information consumer must do if using and onward distributing open energy information. These obligations are typically licencing requirements (i.e., displaying the full text of the data provider's licence, or attribution of share alike conditions).

While this proposed language offers greater sophistication that access/or not, and the ability to create, read, update, and delete, the structure of the rules has limitations. It is also unclear how internal use restrictions would be enforced. For example, if an information user were to use or adopt it for commercial use when the capability granted was strictly for non-commercial use.

5. Open issues

Recent coverage of unauthorised disclosure of sensitive US defence and intelligence data, as well as three major breaches involving UK police personnel information, are indicative of the current problem of controlling access to sensitive information.

5.1 Nature of entity seeking access

In future access control may be required to enforce controls based on location, user identity, system, application, and nature of processing, i.e., an information creator can specify where a user may process the informatoin, on what system(s), using which applications and the permitted processing. Thus, for example, an authorised user may be permitted to access and read a file at several specified locations but only print the file if at a single specific location.

5.2 Identification, authentication, and authorisation

If access control decisions are to consider location, user identity, system, and application, this will require a means to uniquely identify the parameters and to provide an acceptable level of 'identity' authentication. With regards to authorisation, consideration needs to be given to how authorisation is implemented, for example, at individual user/system/application level, at class of user/system/application, and how location is represented (i.e., country, site, building, floor, office, etc.).

5.3 Specification of access control requirements/rules

For access control to be effective the information creator needs to be able to specify and maintain access control requirements or rules for information under their control. As the sophistication of the requirements or rules increases so does the complexity of creating, assuring, and maintaining the control requirements or rules. Consideration needs to be given on how to specify them so that they are readable to machines and people.

5.4 Information model for access control

This issue is associated with the specification of access control requirements or rules. Adoption of a standardised information model for the requirements/rules could ease interoperability between systems.

6. Conclusions and recommendations

Our current models of control are largely based around managing access to systems, applications, folders, and files. This is essentially a binary approach, you either have access or you don't. Whilst there may be some consideration about the level of access at for example at folder level the implication is that the user can access all the folder contents rather than just those items meeting their information needs. Such considerations are generally unlikely to accommodate controls regarding volumetrics, granularity, specific use or purpose to which the access to information maybe put. For sensitive documents, or information containers, this is problematic when the information asset contains significantly more information than is reasonably required for a user to complete a task or make a decision.

The current models lack the granularity and specificity that will be needed in a federated information sharing environment, and consequently there may be significant oversharing of information, creating sensitivity and security issues. To address this, it is recommended that further work is required on the entitlement model. This should be supported by the development of a protocol that allows access and control requirements to be specified in a manner that is comprehensive and can be assured.

References

Atlam, H.F.; Alenezi, A.; Walters, R.J.; Wills, G.B.; Daniel, J. (2017, June) "Developing an adaptive Risk-based access control model for the Internet of Things." In *Proceedings of the 2017 IEEE International Conference on Internet of Things (iThings) and IEEE Green Computing and Communications (GreenCom) and IEEE Cyber, Physical and Social Computing (CPSCom) and IEEE Smart Data (SmartData)*, Exeter, UK, 21–23 June, pp. 655–661.

Atlam, H.F., Alassafi, M.O., Alenezi, A., Walters, R.J. and Wills, G.B. (2018, May) "XACML for Building Access Control Policies in Internet of Things." In *Proceedings of the 3rd International Conference on Internet of Things, Big Data and Security* (IoTBDS 2018), Madeira, Portugal, 19–21 May. pp. 253-260.

Atlam, H.F., Azad, M.A., Alassafi, M.O., Alshdadi, A.A. and Alene, A. (2020) "Risk-Based Access Control Model: A Systematic Literature Review." Future Internet, 12(6), pp.103-126. DOI: 10.3390/fi12060103

Atluri V. and Gal A. (2002) “An authorization model for temporal and derived data: securing information portals.” ACMTrans. Inf. Syst. Secur., 5(1). pp.62–94.

Bertino E., Bettini C., Ferrari E., and Samarati P. (1998) “An access control model supporting periodicity constraints and temporal reasoning.” ACM Trans. Database Syst., 23(3), pp.231–285.

Bertino E., Bonatti P.A., and Ferrari E. (2001) “TRBAC: a temporal role-based access control model.” ACM Trans. Inf. Syst. Secur., 4(3), pp.191–233.

Brewer, D.F.C., and Nash, M.J., (1989) "The Chinese Wall security policy." Proceedings. 1989 IEEE Symposium on Security and Privacy, pp.206-214, DOI: 10.1109/SECPRI.1989.36295.

Cambridge Dictionary - https://dictionary.cambridge.org/dictionary/english/entitlement

Chandran, S.M., and Joshi, J.B.D. (2005). “LoT-RBAC: A Location and Time-Based RBAC Model.” In: Ngu, A.H.H., Kitsuregawa, M., Neuhold, E.J., Chung, JY., Sheng, Q.Z. (eds) Web Information Systems Engineering – WISE 2005. Lecture Notes in Computer Science, vol 3806. Springer, Berlin, Heidelberg. DOI: 10.1007/11581062_27

Copyright, Designs and Patents Act 1988, c.48, Part I, Ch.IV. Available: https://www.legislation.gov.uk/ukpga/1988/48/part/I/chapter/IV

Corradi, A., Montanari, R. and Tibaldi, D. (2004) “Context-Based Access Control Management in Ubiquitous Environments”, Network Computing and Applications, Proceedings of the Third IEEE International Symposium on (NCA'04), August 30 - September 01, 2004, Boston, MA.

Damiani, M.L., Bertino, E., Catania, B. and Perlasca, P., (2007) “GEO-RBAC: a spatially aware RBAC.” ACM Transactions on Information and System Security (TISSEC), 10(1), pp.1-42. DOI: 10.1145/1210263.1210265

DOD (1985) “Trusted Computer System Evaluation Criteria.” DoD 5200.28-STD

Fernandez, E.B., Larrondo-Petrie, M.M. and Escobar, A. E. (2007) "Contexts and Context-Based Access Control," Third International Conference on Wireless and Mobile Communications (ICWMC'07), Guadeloupe, French Caribbean, pp.73-73, DOI: 10.1109/ICWMC.2007.30.

Ferraiolo, D., Kuhn, D and Chandramouli, R. (2003) “Role-Based Access Control”, Artech House.

Fu S. and Xu C.-Z. (2005) “A coordinated spatio-temporal access control model for mobile computing in coalition environments.” In Proc. 19th EEE International Parallel and Distributed Processing Symposium. DOI: 10.1109/IPDPS.2005.10

Gamma, E., Helm, R., Johnson, R. and Vlissides, J. (1994) “Design patterns – Elements of reusable object-oriented software.” Addison-Wesley.

Gov.uk “Holiday entitlement” - https://www.gov.uk/holiday-entitlement-rights

Harrison M. A., Ruzzo W. L., and Ullman J. D., (1976) “Protection in Operating Systems.” Communications of the ACM, vol.19, iss.8, pp.461-471. DOI: 10.1145/360303.360333

Hohfeld, W. N. (1913) “Fundamental Legal Conceptions as Applied in Judicial Reasoning.” The Yale Law Journal, Nov., Vol.23, No.1, pp.16-59.

Joshi J.B.D., Bertino E., Latif U., and Ghafoor A. (2005) “A generalized temporal role-based access control model.” IEEE Trans. Knowl. Data Eng., 17(1), pp.4–23.

Khambhammettu, H.; Boulares, S.; Adi, K.; Logrippo, L. (2013) “A framework for risk assessment in access control systems.” Computers and Security, vol.39, part.A, pp.86–103. DOI: 10.1016/j.cose.2013.03.010

Kirsch-Pinheiro, M., Villanova-Oliver, M., Gensel, J. and Martin, H. (2005, March) “Context-Aware Filtering for Collaborative Web Systems: Adapting the Awareness Information to the User’s Context.” In Procs. of the ACM Symposium on Applied Computing. SAC’05, Santa Fe, New Mexico, USA.

NIST (1994) “Security in Open Systems.” SP 800-7. NIST, Gaithersburg, MD. DOI: 10.6028/NIST.SP.800-7

NIST (2006) “Assessment of Access Control Systems.” NISTIR 7316, NIST, Gaithersburg, MD. DOI: 10.6028/NIST.IR.7316

NIST (2014) “Guide to Attribute Based Access Control (ABAC) Definition and Considerations.” SP 800-162. NIST, Gaithersburg, MD. DOI: 10.6028/NIST.SP.800-162

Oh, S. and Park, S., (2003) “Task–role-based access control model.” Information systems, vol.28, iss.6, pp.533-562. DOI: 10.1016/S0306-4379(02)00029-7

Open Energy (2022a) “Access Control and Capability Grant Language.“ Icebreaker One Limited, v:1.0.0. Available: https://docs.openenergy.org.uk/1.0.0/access_control_specification.html

Open Energy (2022b) “Data Access Conditions.“ Icebreaker One Limited, v:1.0.0. Available: https://docs.openenergy.org.uk/1.0.0/ops_guidelines/common_policies.html#data-access-conditions

Ray I., Kumar M., and Yu L. (2006, December) “LRBAC: A location-aware role-based access control model.” In Information Systems Security: Second International Conference, ICISS 2006, Kolkata, India. Proceedings 2, pp. 147-161. Springer Berlin Heidelberg. DOI: 10.1007/11961635_10

Ray, I., Toahchoodee, M. (2007) “A Spatio-temporal Role-Based Access Control Model. In: Barker, S., Ahn, GJ. (eds) Data and Applications Security XXI. DBSec 2007. Lecture Notes in Computer Science, vol 4602. Springer, Berlin, Heidelberg. DOI: 10.1007/978-3-540-73538-0_16

Risse, M. (2023) “Political Theory of the Digital Age: Where Artificial Intelligence Might Take Us.” Cambridge University Press, Cambridge. DOI: 10.1017/9781009255189

Sandhu R.S., and Samarati, P. (September 1994) “Access Control: Principles and Practice.” IEEE Communications Magazine, vol.32, iss.9, pp.40-48. DOI: 10.1109/35.312842

Shue, H. (1980). “Basic Rights: Subsistence, Affluence, and U.S. Foreign Policy.” Princeton University Press. Chapters 1-2.